Max-Weight Revisited: Sequences of Non-Convex Optimisations Solving Convex Optimisations

Víctor Valls, Douglas J. Leith
Trinity College Dublin

arXiv:1406.0899v4 [math.OC] 26 Aug 2015

***Abstract*—We investigate the connections between max-weight approaches and dual subgradient methods for convex optimisation. We find that strong connections exist and we establish a clean, unifying theoretical framework that includes both max-weight and dual subgradient approaches as special cases. Our analysis uses only elementary methods, and is not asymptotic in nature. It also allows us to establish an explicit and direct connection between discrete queue occupancies and Lagrange multipliers.**

***Index Terms*—convex optimisation, max-weight scheduling, backpressure, subgradient methods.**

I. Introduction

In queueing networks, max-weight (also referred to as backpressure) approaches have been the subject of much interest for solving utility optimisation problems in a distributed manner.

In brief, consider a queueing network where the queue occupancy of the i'th queue at time k is denoted by $Q_k^{(i)} \in \mathbb{N}$, $i = 1, 2, \ldots, n$, and we gather these together into vector $\boldsymbol{Q}_k \in \mathbb{N}^n$. Time is slotted and at each time step $k = 1, 2, \ldots$ we select action $\boldsymbol{x}_k \in D \subset \mathbb{N}^n$, *e.g.*, selecting i'th element $x_k^{(i)} = 1$ corresponds to transmitting one packet from queue i and $x_k^{(i)} = 0$ to doing nothing. The connectivity between queues is captured via matrix $\boldsymbol{A} \in \{-1, 0, 1\}^{n \times n}$, whose i'th row has a -1 at the i'th entry, 1 at entries corresponding to queues from which packets are sent to queue i, and 0 entries elsewhere. The queue occupancy then updates according to $\boldsymbol{Q}_{k+1} = [\boldsymbol{Q}_k + \boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{b}_k]^+$, $i = 1, 2, \ldots, n$, where the i'th element of vector $\boldsymbol{b}_k \in \mathbb{N}^n$ denotes the number of external packet arrivals to queue i at time k. The objective is to stabilise all of the queues while maximising utility $U(\boldsymbol{z}_k)$ where $U : \mathbb{R}^n \to \mathbb{R}$ is concave and continuously differentiable and $\boldsymbol{z}_k$ is a running average of $\boldsymbol{x}_j$, $j = 1, \ldots, k$. The *greedy primal-dual* variant of max-weight scheduling [1], for example, selects action $\boldsymbol{x}_k \in \arg\max_{\boldsymbol{x} \in D} \partial U(\boldsymbol{z}_k)^T \boldsymbol{x} - \beta \boldsymbol{Q}_k^T \boldsymbol{A}\boldsymbol{x}$ with $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta \boldsymbol{x}_k$, $0 < \beta < 1$ a design parameter.

Appealing features of this max-weight scheduling approach include the lack of a requirement for *a priori* knowledge of packet arrival process $\{\boldsymbol{b}_k\}$, and the fact that the discrete action set matches the actual decision variables (namely, do we transmit a packet or not). Importantly, although cost function $-U(\cdot)$ is required to be convex, at each time step the max-weight optimisation is non-convex owing to the non-convexity of action set D. Further, convergence is typically proved using Foster-Lyapunov or by sophisticated fluid-limit arguments, which allow sequence $\{\boldsymbol{b}_k\}$ to be accommodated but are distinct from the usual approaches employed in convex optimisation. Hence, the body of work on max-weight approaches remains separate from the mainstream literature on convex optimisation. On the other hand, queueing and Lagrange multiplier subgradient updates are clearly similar, at least superficially, although the exact nature of the relationship between queues and multipliers remains unclear.

Taking these observations as our starting point, in this paper we investigate the connections between max-weight approaches and dual subgradient methods for convex optimisation. We find that strong connections do indeed exist and we establish a clean, unifying theoretical framework that includes both max-weight and dual subgradient approaches as special cases. In summary, the main contributions of the paper include the following.

1) *Generalising max-weight.* Our analysis places max-weight firmly within the field of convex optimisation, extending it from the specific constraints induced by queueing networks to general convex nonlinear contraints with bounded curvature. We show that any non-convex update with suitable descent properties can be employed, and the wealth of convex descent methods can be leveraged to derive non-convex approaches. Descent methods studied here include non-convex variants of the classical Frank-Wolfe update and of the primal Lagrangian update.

2) *Generalising dual subgradient methods.* We show that convexity can be relaxed in classical dual subgradient methods, allowing use of a finite action set. In the special case of optimisation problems with linear constraints, we rigorously establish a close connection (essentially an equivalence) between Lagrange multiplier subgradient updates and discrete queues, so putting existing intuition on a sound footing.

3) *Unifying theoretical framework.* In generalising max-weight and dual subgradient methods our analysis clarifies the fundamental properties required. In particular, bounded curvature of the objective and constraint functions plays a prominent role in our analysis, as does boundedness of the action set. Of interest in its own right, we note that our analysis requires only elementary methods and so an additional contribution is the accessible nature of the methods of proof employed. In particular, it turns out that deterministic analysis of sample paths is sufficient to handle stochasticity. The methods of proof themselves are new in the context of max-weight approaches, and are neither Foster-Lyapunov nor fluid-limit based.

This work was supported by Science Foundation Ireland under Grant No. 11/PI/1177.

A. Related Work

Max-weight scheduling was introduced by Tassiulas and Ephremides in their seminal paper [2]. They consider a network of queues with slotted time, an integer number of packet arrivals in each slot and a finite set of admissible scheduling patterns, referred to as *actions*, in each slot. Using a Forster-Lyapunov approach they present a scheduling policy that stabilises the queues provided the external traffic arrivals are strictly feasible. Namely, the scheduling policy consists of selecting the action at each slot that maximises the queue-length-weighted sum of rates, $\boldsymbol{x}_k \in \arg\max_{\boldsymbol{x}\in D} -\boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x}$.

Independently, [1], [3], [4] proposed extensions to the max-weight approach to accommodate concave utility functions. In [1] the *greedy primal-dual* algorithm is introduced, as already described above, for network linear constraints and utility function $U(\cdot)$ which is continuously differentiable and concave. The previous work is extended in [5] to consider general nonlinear constraints. In [4] the utility fair allocation of throughput in a cellular downlink is considered. The utility function is of the form $U(\boldsymbol{z}) = \sum_{i=1}^n U_i(z^{(i)})$, $U_i(z) = \beta_i(z^{(1-\frac{1}{m})})/(1-\frac{1}{m})$, with m, β_i design parameters. Queue departures are scheduled according to $\boldsymbol{x}_k \in \arg\max_{\boldsymbol{x}\in\mathrm{conv}(D)} -\boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x}$ and queue arrivals are scheduled by a congestion controller such that $E[b_k^{(i)}|\boldsymbol{Q}_k] = \min\{\partial U_i(Q_k^{(i)}), M\}$ and $E[(b_k^{(i)})^2|\boldsymbol{Q}_k] \le A$ where A, M are positive constants. The work in [3] considers power allocation in a multibeam downlink satellite communication link with the aim of maximising throughput while ensuring queue stability. This is extended in a sequence of papers [6], [7], [8] and a book [9] to develop the *drift plus penalty* approach. In this approach the basic strategy for scheduling queue departures is according to $\boldsymbol{x}_k \in \arg\max_{\boldsymbol{x}\in D} -\boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x}$ and utility functions are incorporated in a variety of ways. For example, for concave non-decreasing continuous utility functions U of the form $U(\boldsymbol{z}) = \sum_{i=1}^n U_i(z^{(i)})$ one formulation is for a congestion controller to schedule arrivals into an ingress queue such that $b_k^{(i)} \in \arg\max_{0\le b\le R} VU_i(b) - bQ_k^{(i)}$ where V, R are sufficiently large design parameters and $b \in \mathbb{R}$ [10]. Another example is for cost functions of the form $E[P_k(\boldsymbol{x}_k)]$ where $P_k(\cdot)$ is bounded, i.i.d. and known at each time step, in which case actions at each time step are selected to minimise $\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} VP_k(\boldsymbol{x}_k) + \boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x}$ where V is a design parameter [9].

With regard to the existence of a connection between the discrete-valued queue occupancy in a queueing network and continuous-valued Lagrange multipliers, this has been noted by several authors, see for example [11], [12], and so might be considered something of a "folk theorem" but we are aware of few rigorous results. A notable exception is [13], which establishes that a discrete queue update tends on average to drift towards the optimal multiplier value. Also, the *greedy primal-dual* algorithm presented in [1] shows that asymptotically as design parameter $\beta \to 0$ and $t \to \infty$ the scaled queue occupancy converges to the set of dual optima.

Selection of a sequence of actions in a discrete-like manner is also considered in the convex optimisation literature. The *nonlinear Gauss-Seidel* algorithm, also known as *block coordinate descent* [14], [15] minimises a convex function over a convex set by updating one co-ordinate at a time. The convex function is required to be continuously differentiable and strictly convex and, unlike in the max-weight algorithms discussed above, the action set is convex. The classical Frank-Wolfe algorithm [16] also minimises a convex continuously differentiable function over a polytope by selecting from a discrete set of descent directions, although a continuous-valued line search is used to determine the final update. We also note the work on online convex optimisation [17], [18], where the task is to choose a sequence of actions so to minimise an unknown sequence of convex functions with low regret.

B. Notation

Vectors and matrices are indicated in bold type. Since we often use subscripts to indicate elements in a sequence, to avoid confusion we usually use a superscript $x^{(j)}$ to denote the j'th element of a vector $\boldsymbol{x}$. The j'th element of operator $[\boldsymbol{x}]^{[0,\bar\lambda]}$ equals $x^{(j)}$ (the j'th element of $\boldsymbol{x}$) when $x^{(j)} \in [0,\bar\lambda]$ and otherwise equals 0 when $x^{(j)} < 0$ and $\bar\lambda$ when $x^{(j)} > \bar\lambda$. Note that we allow $\bar\lambda = +\infty$, and following standard notation in this case usually write $[x]^+$ instead of $[x]^{[0,\infty)}$. The subgradient of a convex function f at point $\boldsymbol{x}$ is denoted $\partial f(\boldsymbol{x})$. For two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$ we use element-wise comparisons $\boldsymbol{x} \succeq \boldsymbol{y}$ and $\boldsymbol{y} \succ \boldsymbol{x}$ to denote when $y^{(j)} \ge x^{(j)}$, $y^{(j)} > x^{(j)}$ respectively for all $j = 1, \ldots, m$.

II. Preliminaries

We recall the following convexity properties.

Lemma 1 (Lipschitz Continuity)**.** *Let $h : M \to \mathbb{R}$ be a convex function and let C be a closed and bounded set contained in the relative interior of the domain $M \subseteq \mathbb{R}^n$. Then $h(\cdot)$ is Lipschitz continuous on C* i.e.*, there exists constant ν_h such that $|h(\boldsymbol{x}) - h(\boldsymbol{y})| \le \nu_h\|\boldsymbol{x}-\boldsymbol{y}\|_2\ \forall \boldsymbol{x}, \boldsymbol{y} \in C$.*

Proof: See, for example, [19]. ■

Lemma 2 (Bounded Distance)**.** *Let $D := \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{|D|}\}$ be a finite set of points from $\mathbb{R}^n$. Then there exists constant $\bar x_D$ such that $\|\boldsymbol{z}-\boldsymbol{y}\|_2 \le \bar x_D$ for any two points $\boldsymbol{z}, \boldsymbol{y} \in C := \mathrm{conv}(D)$, where $\mathrm{conv}(D)$ denotes the convex hull of D.*

Proof: Since $\boldsymbol{z}, \boldsymbol{y} \in C$ these can be written as the convex combination of points in D, *i.e.*, $\boldsymbol{z} = \sum_{j=1}^{|D|} a^{(j)}\boldsymbol{x}_j$, $\boldsymbol{y} = \sum_{j=1}^{|D|} b^{(j)}\boldsymbol{x}_j$ with $\|\boldsymbol{a}\|_1 = 1 = \|\boldsymbol{b}\|_1$. Hence $\|\boldsymbol{z}-\boldsymbol{y}\|_2 = \|\sum_{j=1}^{|D|}(a^{(j)} - b^{(j)})\boldsymbol{x}_j\|_2 \le \sum_{j=1}^{|D|}\|a^{(j)} - b^{(j)}\|_2\|\boldsymbol{x}_j\|_2 \le \bar x_D := 2\max_{\boldsymbol{x}\in D}\|\boldsymbol{x}\|_2$. ■

We also introduce the following definition:

Definition 1 (Bounded Curvature)**.** *Let $h : M \to \mathbb{R}$ be a convex function defined on domain $M \subseteq \mathbb{R}^n$. We say the $h(\cdot)$ has bounded curvature on set $C \subset M$ if for any points $\boldsymbol{z}, \boldsymbol{z}+\boldsymbol{\delta} \in C$*

$$h(\boldsymbol{z}+\boldsymbol{\delta}) - h(\boldsymbol{z}) \le \partial h(\boldsymbol{z})^T\boldsymbol{\delta} + \mu_h\|\boldsymbol{\delta}\|_2^2 \qquad (1)$$

where $\mu_h \ge 0$ is a constant that does not depend on $\boldsymbol{z}$ or $\boldsymbol{\delta}$.

Bounded curvature will prove important in our analysis. The following lemma shows that a necessary and sufficient

condition for bounded curvature is that the subgradients of $h(\cdot)$ are Lipschitz continuous on set C.

Lemma 3 (Bounded Curvature)**.** *Let $h : M \to \mathbb{R}$, $M \subseteq \mathbb{R}^n$ be a convex function. Then $h(\cdot)$ has bounded curvature on C if and only if for all $\boldsymbol{z}, \boldsymbol{z}+\boldsymbol{\delta} \in C$ there exists a member $\partial h(\boldsymbol{z})$ (respectively, $\partial h(\boldsymbol{z}+\boldsymbol{\delta})$) of the set of subdifferentials at point $\boldsymbol{z}$ (respectively, $\boldsymbol{z}+\boldsymbol{\delta}$) such that $(\partial h(\boldsymbol{z}+\boldsymbol{\delta}) - \partial h(\boldsymbol{z}))^T \boldsymbol{\delta} \leq \mu_h \|\boldsymbol{\delta}\|_2^2$ where μ_h does not depend on $\boldsymbol{z}$ or $\boldsymbol{\delta}$.*

Proof: $\Rightarrow$ Suppose $h(\cdot)$ has bounded curvature on C. From (1) it follows that $h(\boldsymbol{z}+\boldsymbol{\delta}) - h(\boldsymbol{z}) \leq \partial h(\boldsymbol{z})^T\boldsymbol{\delta} + \mu_h\|\boldsymbol{\delta}\|_2^2$ and $h(\boldsymbol{z}) - h(\boldsymbol{z}+\boldsymbol{\delta}) \leq -\partial h(\boldsymbol{z}+\boldsymbol{\delta})^T\boldsymbol{\delta} + \mu_h\|\boldsymbol{\delta}\|_2^2$. Adding left-hand and right-hand sides of these inequalities yields $0 \leq (\partial h(\boldsymbol{z}) - \partial h(\boldsymbol{z}+\boldsymbol{\delta}))^T \boldsymbol{\delta} + 2\mu_h\|\boldsymbol{\delta}\|_2^2$ *i.e.,* $(\partial h(\boldsymbol{z}+\boldsymbol{\delta}) - \partial h(\boldsymbol{z}))^T \boldsymbol{\delta} \leq \mu_h\|\boldsymbol{\delta}\|_2^2$.

$\Leftarrow$ Suppose $(\partial h(\boldsymbol{z}+\boldsymbol{\delta}) - \partial h(\boldsymbol{z}))^T \boldsymbol{\delta} \leq \mu_h\|\boldsymbol{\delta}\|_2$ for all $\boldsymbol{z}, \boldsymbol{z}+\boldsymbol{\delta} \in M$. It follows that $\partial h(\boldsymbol{z}+\boldsymbol{\delta})^T\boldsymbol{\delta} \leq \partial h(\boldsymbol{z})^T\boldsymbol{\delta} + \mu_h\|\boldsymbol{\delta}\|_2^2$. By the definition of the subgradient we have that $h(\boldsymbol{z}+\boldsymbol{\delta}) - h(\boldsymbol{z}) \leq \partial h(\boldsymbol{z}+\boldsymbol{\delta})^T\boldsymbol{\delta}$, and so we obtain that $h(\boldsymbol{z}+\boldsymbol{\delta}) - h(\boldsymbol{z}) \leq \partial h(\boldsymbol{z})^T\boldsymbol{\delta} + \mu_h\|\boldsymbol{\delta}\|_2^2$. ■

III. Non-Convex Descent

We begin by considering minimisation of convex function $F : \mathbb{R}^n \to \mathbb{R}$ on convex set $C := \text{conv}(D)$, the convex hull of set $D := \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{|D|}\}$ consisting of a finite collection of points from $\mathbb{R}^n$ (so C is a polytope). Our interest is in selecting a sequence of points $\{\boldsymbol{x}_k\}$, $k = 1, 2, \ldots$ from set D such that the running average $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$ minimises $F(\cdot)$ for k sufficiently large and β sufficiently small. Note that set D is non-convex since it consists of a finite number of points, and by analogy with max-weight terminology we will refer to it as the *action set*.

Since C is the convex hull of action set D, any point $\boldsymbol{z}^* \in C$ minimising $F(\cdot)$ can be written as convex combinations of points in D *i.e.,* $\boldsymbol{z}^* = \sum_{j=1}^{|D|} \theta^{*(j)}\boldsymbol{x}_j$, $\theta^{*(j)} \in [0,1]$, $\|\boldsymbol{\theta}\|_1 = 1$. Hence, we can always construct sequence $\{\boldsymbol{x}_k\}$ by selecting points from set D in proportion to the $\theta^{*(j)}$, $j = 1, \ldots, |D|$. That is, by *a posteriori* time-sharing (*a posteriori* in the sense that we need to find minimum $\boldsymbol{z}^*$ before we can construct sequence $\{\boldsymbol{x}_k\}$). Of more interest, however, it turns out that when function $F(\cdot)$ has bounded curvature then sequences $\{\boldsymbol{x}_k\}$ can be found without requiring knowledge of $\boldsymbol{z}^*$.

A. Non-Convex Direct Descent

The following theorem formalises the above commentary, also generalising it to sequences of convex functions $\{F_k\}$ rather than just a single function as this will prove useful later.

Theorem 1 (Greedy Non-Convex Convergence)**.** *Let $\{F_k\}$ be a sequence of convex functions with uniformly bounded curvature μ_F on set $C := \text{conv}(D)$, action set D a finite set of points from $\mathbb{R}^n$. Let $\{\boldsymbol{z}_k\}$ be a sequence of vectors satisfying $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$ with $\boldsymbol{z}_1 \in C$ and*

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} F_k((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}), \;\; k = 1, 2, \ldots \quad (2)$$

Suppose parameter β is sufficiently small that

$$0 < \beta \leq (1-\gamma)\min\{\epsilon/(\mu_F \bar{x}_D^2), 1\} \quad (3)$$

with $\epsilon > 0$, $\gamma \in (0,1)$, $\bar{x}_D := 2\max_{\boldsymbol{x}\in D}\|\boldsymbol{x}\|_2$ and that functions F_k change sufficiently slowly that

$$|F_{k+1}(\boldsymbol{z}) - F_k(\boldsymbol{z})| \leq \gamma_1\gamma\beta\epsilon, \;\; \forall \boldsymbol{z} \in C$$

with $\gamma_1 \in (0, 1/2)$. Then for every $\epsilon > 0$ and k sufficiently large we have that

$$0 \leq F_k(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{y}_k^*) \leq 2\epsilon$$

where $\boldsymbol{y}_k^* \in \arg\min_{\boldsymbol{z}\in C} F_k(\boldsymbol{z})$.

Proof: See Appendix. ■

Observe that in Theorem 1 we select $\boldsymbol{x}_k$ by solving non-convex optimisation (2) at each time step. This optimisation is one step ahead, or greedy, in nature and does not look ahead to future values of the sequence or require knowledge of optima $\boldsymbol{y}_k^*$. Of course, such an approach is mainly of interest when non-convex optimisation (2) can be efficiently solved, *e.g.,* when action set D is small or the optimisation separable.

Observe also that Theorem 1 relies upon the bounded curvature of the sequence of functions $F_k(\cdot)$. A smoothness assumption of this sort seems essential, since when it does not hold it is easy to construct examples where Theorem 1 does not hold. Such an example is illustrated schematically in Figure 1a. The shaded region in Figure 1a indicates the level set $\{F(\boldsymbol{y}) \leq F(\boldsymbol{z}) : \boldsymbol{y} \in C\}$. The level set is convex, but the boundary is non-smooth and contains "kinks". We can select points from the set $\{(1-\beta)\boldsymbol{z} + \beta\boldsymbol{x} : \boldsymbol{x} \in D = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3\}\}$. This set of points is indicated in Figure 1a and it can be seen that every point lies outside the level set. Hence, we must have $F((1-\beta)\boldsymbol{z} + \beta\boldsymbol{x}) > F(\boldsymbol{z})$, and upon iterating we will end up with a diverging sequence. Note that in this example changing the step size β does not resolve the issue. Bounded curvature ensures that the boundary of the level sets is smooth, and this ensures that for sufficiently small β there exists a convex combination of $\boldsymbol{z}$ with a point $\boldsymbol{x} \in D$ such that $F((1-\beta)\boldsymbol{z} + \beta\boldsymbol{x}) < F(\boldsymbol{z})$ and so the solution to optimisation (2) improves our objective, see Figure 1b.

Theorem 1 is stated in a fairly general manner since this will be needed for our later analysis. An immediate corollary to Theorem 1 is the following convergence result for unconstrained optimisation.

Corollary 1 (Unconstrained Optimisation)**.** *Consider the following sequence of non-convex optimisations $\{P_k^u\}$:*

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} f\left((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}\right)$$
$$\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$$

with $\boldsymbol{z}_1 \in C := \text{conv}(D)$, action set $D \subset \mathbb{R}^n$ finite. Then $0 \leq f(\boldsymbol{z}_k) - f^* \leq 2\epsilon$ for all k sufficiently large, where $f^* = \min_{\boldsymbol{z}\in C} f(\boldsymbol{z})$, provided $f(\cdot)$ has bounded curvature with curvature constant μ_f and $0 < \beta \leq (1-\gamma)\min\{\epsilon/(\mu_f\bar{x}_D^2), 1\}$ with $\gamma \in (0,1)$, $\epsilon > 0$, $\bar{x}_D := 2\max_{\boldsymbol{x}\in D}\|\boldsymbol{x}\|_2$.

Figure 2 illustrates Corollary 1 schematically in $\mathbb{R}^2$. The sequence of non-convex optimisations descends in two iterations $f(\boldsymbol{z}_1) > f(\boldsymbol{z}_2) > f(\boldsymbol{z}_3)$ (using points $\boldsymbol{x}_3$ and $\boldsymbol{x}_4$ respectively) and $f(\boldsymbol{z}_k) - f^* \leq 2\epsilon$ for $k > 3$ (not shown in Figure 2).

Note that the curvature constant μ_f of function f need not be known, an upper bound being sufficient to select β. Next

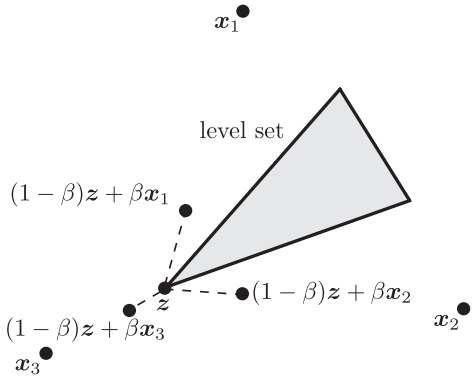

(a) Example where $F(\cdot)$ does not possess bounded curvature. None of $(1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}$, $\boldsymbol{x} \in D$ are descending.

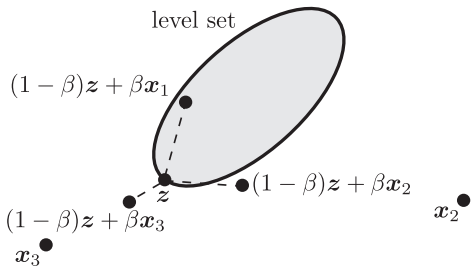

(b) Example where $F(\cdot)$ has bounded curvature. For β sufficiently small, for at least one $(1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}$, $\boldsymbol{x} \in D$ descent is possible.

Fig. 1: Illustrating how bounded curvature allows monotonic descent. Set D consists of the marked points $\boldsymbol{x}_1$, $\boldsymbol{x}_2$, $\boldsymbol{x}_3$. Level set $\{F(\boldsymbol{y}) \leq F(\boldsymbol{z}) : \boldsymbol{y} \in C\}$ is indicated by the shaded areas. The possible choices of $(1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}$, $\boldsymbol{x} \in D$ are indicated.

we present two brief examples that are affected differently by constant μ_f.

Example 1 (Linear Objective)**.** Suppose $f(\boldsymbol{z}) := \boldsymbol{a}^T\boldsymbol{z}$ where $\boldsymbol{a} \in \mathbb{R}^n$. The objective function is linear and so has curvature constant $\mu_f = 0$. It can be seen from (3) that we can choose β independently of parameter ϵ. Further, for any $\beta \in (0,1)$ we have that $f(\boldsymbol{z}_{k+1}) < f(\boldsymbol{z}_k)$ for all k and so $f(\boldsymbol{z}_k) \to f^*$.

Example 2 (Quadratic Objective)**.** Suppose $f(\boldsymbol{z}) := \frac{1}{2}\boldsymbol{z}^T\boldsymbol{A}\boldsymbol{z}$ where $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ is positive definit. Then $\mu_f = \lambda_{\max}(\boldsymbol{A}) > 0$ and in contrast to Example 1 the bound (3) on parameter β now depends on ϵ and convergence is into the ball $f(\boldsymbol{z}_k) - f^* \leq 2\epsilon$ for k sufficiently large.

B. Non-Convex Frank-Wolfe-like Descent

It is important to note that other convergent non-convex updates are also possible. For example:

Theorem 2 (Greedy Non-Convex FW Convergence)**.** *Consider the setup in Theorem 1, but with modified update*

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{x},\ k = 1, 2, \ldots \tag{4}$$

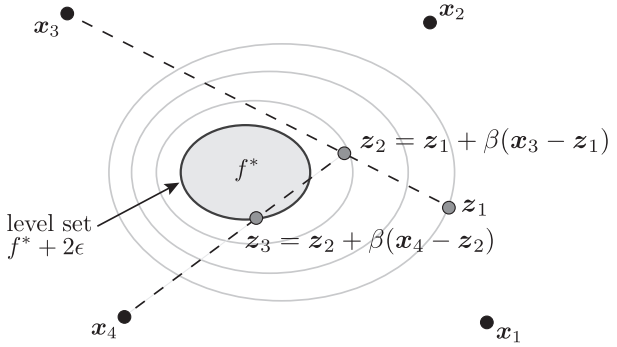

Fig. 2: Illustrating unconstrained convergence in $\mathbb{R}^2$. The sequence of non-convex optimisations converges with $k = 2$. The function average decreases monotonically and then remains in level set $f(\boldsymbol{z}_k) \leq f^* + 2\epsilon$ for $k \geq 3$.

Then for every $\epsilon > 0$ and k sufficiently large we have that

$$0 \leq F_k(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{y}_k^*) \leq 2\epsilon$$

where $\boldsymbol{y}_k^* \in \arg\min_{\boldsymbol{z}\in C} F_k(\boldsymbol{z})$.

Proof: See Appendix. ∎

The intuition behind the update in Theorem 2 is that at each step we locally approximate $F_k(\boldsymbol{z}_{k+1})$ by linear function $F_k(\boldsymbol{z}_k) + \partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{z}_{k+1} - \boldsymbol{z}_k)$ and then minimise this linear function. Since $F_k(\cdot)$ is convex, this linear function is in fact the supporting hyperplane to $F_k(\cdot)$ at point $\boldsymbol{z}_k$, and so can be expected to allow us to find a descent direction. Similar intuition also underlies classical Frank-Wolfe algorithms for convex optimisation [16] on a polytope, and Theorem 2 extends this class of algorithms to make use of non-convex update (4) and a fixed step size (rather than the classical approach of selecting the step size by line search).

Note that when the function is linear $F_k(\boldsymbol{z}) = \boldsymbol{c}_k^T\boldsymbol{z}$, $\boldsymbol{c}_k \in \mathbb{R}^n$, then $\arg\min_{\boldsymbol{x}\in D} F_k((1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}) = \arg\min_{\boldsymbol{x}\in D} \boldsymbol{c}_k^T\boldsymbol{x}$ and $\arg\min_{\boldsymbol{x}\in D} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{x} = \arg\min_{\boldsymbol{x}\in D} \boldsymbol{c}_k^T\boldsymbol{x}$. That is, updates (2) and (4) are identical.

Note also that

$$\arg\min_{\boldsymbol{x}\in D} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{x} \subseteq \arg\min_{\boldsymbol{z}\in C} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{z}. \tag{5}$$

This is because the RHS of (5) is a linear program (the objective is linear and set C is a polytope, so defined by linear constraints) and so the optimum set is either (i) an extreme point of C and so a member of set D, or (ii) a face of polytope C with the extreme points of the face belonging to set D. Hence, while update (4) is non-convex it can nevertheless be solved in polynomial time.

IV. Sequences of Non-Convex Optimisations & Constrained Convex Optimisation

We now extend consideration to the constrained convex optimisation P:

$$\begin{aligned} \underset{\boldsymbol{z}\in C}{\text{minimise}} \quad & f(\boldsymbol{z}) \\ \text{subject to} \quad & \boldsymbol{g}(\boldsymbol{z}) \preceq 0 \end{aligned}$$

where $\boldsymbol{g}(\boldsymbol{z}) := [g^{(1)}, \ldots, g^{(m)}]^T$ and $f, g^{(j)} : \mathbb{R}^n \to \mathbb{R}$, $j = 1, \ldots, m$ are convex functions with bounded curvature with, respectively, curvature constants μ_f and $\mu_{g^{(j)}}$. As before, action set D consists of a finite set of points in $\mathbb{R}^n$ and $C := \text{conv}(D)$. Let $C_0 := \{\boldsymbol{z} \in C \mid \boldsymbol{g}(\boldsymbol{z}) \preceq \boldsymbol{0}\}$ denote the set of feasible points, which we will assume has non-empty relative interior (*i.e.,* a Slater point exists). Let $C^* := \arg\min_{\boldsymbol{z} \in C_0} f(\boldsymbol{z}) \subseteq C_0$ be the set of optima and $f^* := f(\boldsymbol{z}^*)$, $\boldsymbol{z}^* \in C^*$.

In the next sections we introduce a generalised dual subgradient approach for finding approximate solutions to optimisation P which, as we will see, includes the classical convex dual subgradient method as a special case.

A. Lagrangian Penalty

As in classical convex optimisation we define Lagrangian $L(\boldsymbol{z}, \boldsymbol{\lambda}) := f(\boldsymbol{z}) + \boldsymbol{\lambda}^T \boldsymbol{g}(\boldsymbol{z})$ where $\boldsymbol{\lambda} = [\lambda^{(1)}, \ldots, \lambda^{(m)}]^T$ with $\lambda^{(j)} \geq 0$, $j = 1, \ldots, m$. Since set C_0 has non-empty relative interior, the Slater condition is satisfied and strong duality holds. That is, there is zero duality gap and so the solution of the dual problem P^D:

$$\underset{\boldsymbol{\lambda} \succeq \boldsymbol{0}}{\text{maximise}} \quad q(\boldsymbol{\lambda}) := \min_{\boldsymbol{z} \in C} L(\boldsymbol{z}, \boldsymbol{\lambda})$$

and primal problem P coincide. Therefore, we have that

$$\min_{\boldsymbol{z} \in C} \max_{\boldsymbol{\lambda} \succeq 0} L(\boldsymbol{z}, \boldsymbol{\lambda}) = \max_{\boldsymbol{\lambda} \succeq 0} \min_{\boldsymbol{z} \in C} L(\boldsymbol{z}, \boldsymbol{\lambda}) = q(\boldsymbol{\lambda}^*) = f^*$$

where $\boldsymbol{\lambda}^* := \arg\max_{\boldsymbol{\lambda} \succeq \boldsymbol{0}} q(\boldsymbol{\lambda})$.

1) Lagrangian Bounded Curvature: As already noted, bounded curvature plays a key role in ensuring convergence to an optimum when selecting from a discrete set of actions. For any two points $\boldsymbol{z}, \boldsymbol{z} + \boldsymbol{\delta} \in C$ we have that

$$L(\boldsymbol{z} + \boldsymbol{\delta}, \boldsymbol{\lambda}) \leq L(\boldsymbol{z}, \boldsymbol{\lambda}) + \partial_{\boldsymbol{z}} L(\boldsymbol{z}, \boldsymbol{\lambda})^T \boldsymbol{\delta} + \mu_L \|\boldsymbol{\delta}\|_2^2,$$

where $\mu_L = \mu_f + \boldsymbol{\lambda}^T \boldsymbol{\mu}_{\boldsymbol{g}}$ with $\boldsymbol{\mu}_{\boldsymbol{g}} := [\mu_{g^{(1)}}, \ldots, \mu_{g^{(m)}}]^T$. It can be seen that the curvature constant μ_L of the Lagrangian depends on the multiplier $\boldsymbol{\lambda}$. Since set $\boldsymbol{\lambda} \succeq \boldsymbol{0}$ is unbounded, it follows that the Lagrangian does not have bounded curvature on this set unless $\boldsymbol{\mu}_{\boldsymbol{g}} = \boldsymbol{0}$ (corresponding to the special case where the constraints are linear). Fortunately, by constraining $\lambda^{(j)} \leq \bar{\lambda}$, $j = 1, \ldots, m$ for some $\bar{\lambda} \geq 0$ resolves the issue, *i.e.,* now $L(\cdot, \boldsymbol{\lambda})$ has uniform bounded curvature with constant

$$\bar{\mu}_L = \mu_f + \bar{\lambda} \boldsymbol{1}^T \boldsymbol{\mu}_{\boldsymbol{g}}.$$

For bounded curvature we only require constant $\bar{\lambda}$ to be finite, but as we will see later in Lemmas 5 and 7 in general it should be chosen with some care.

B. Non-Convex Dual Subgradient Update

In this section we present a primal-dual-like approach in which we use discrete actions to obtain approximate solutions to problem P. In particular, we construct a sequence $\{\boldsymbol{z}_k\}$ of points in C such that $f(\frac{1}{k}\sum_{i=1}^k \boldsymbol{z}_{i+1})$ is arbitrarily close to f^* for k sufficiently large.

We start by introducing two lemmas, which will play a prominent role in later proofs.

Lemma 4 (Minimising Sequence of Lagrangians)**.** *Let $\{\boldsymbol{\lambda}_k\}$ be a sequence of vectors in $\mathbb{R}^m$ such that $\boldsymbol{\lambda}_k \preceq \bar{\lambda}\boldsymbol{1}$, $\bar{\lambda} > 0$ and $\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|_2 \leq \gamma_1 \gamma \beta \epsilon / (m\bar{g})$ with $\gamma \in (0, 1)$, $\gamma_1 \in (0, 1/2)$, β, $\epsilon > 0$, $\bar{g} := \max_{\boldsymbol{z} \in C} \|\boldsymbol{g}(\boldsymbol{z})\|_\infty$. Consider optimisation problem P and updates*

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x} \in D} L((1-\beta)\boldsymbol{z}_k + \beta \boldsymbol{x}, \boldsymbol{\lambda}_k), \tag{6}$$

$$\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta \boldsymbol{x}_k. \tag{7}$$

Then, for k sufficiently large ($k \geq \bar{k}$) we have that $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - q(\boldsymbol{\lambda}_k) \leq L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - f^* \leq 2\epsilon$ provided β is sufficiently small, i.e., *$0 < \beta \leq (1-\gamma)\min\{\epsilon/(\bar{\mu}_L \bar{x}_D^2), 1\}$ where $\bar{x}_D := 2\max_{\boldsymbol{x} \in D} \|\boldsymbol{x}\|_2$, $\bar{\mu}_L = \mu_f + \bar{\lambda}\boldsymbol{1}^T \boldsymbol{\mu}_{\boldsymbol{g}}$.*

Proof: Observe that since $|L(\boldsymbol{z}, \boldsymbol{\lambda}_{k+1}) - L(\boldsymbol{z}, \boldsymbol{\lambda}_k)| = |(\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k)^T \boldsymbol{g}(\boldsymbol{z})| \leq \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|_2 \|\boldsymbol{g}(\boldsymbol{z})\|_2 \leq \|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\lambda}_k\|_2 m\bar{g} \leq \gamma_1 \gamma \beta \epsilon$ and $L(\cdot, \boldsymbol{\lambda}_k)$ has uniformly bounded curvature by Theorem 1 we have that for k sufficiently large ($k \geq \bar{k}$) then $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - q(\boldsymbol{\lambda}_k) \leq 2\epsilon$ where $q(\boldsymbol{\lambda}) := \min_{\boldsymbol{z} \in C} L(\boldsymbol{z}, \boldsymbol{\lambda})$. Further, since $q(\boldsymbol{\lambda}) \leq q(\boldsymbol{\lambda}^*) \leq f^*$ for all $\boldsymbol{\lambda} \succeq \boldsymbol{0}$ it follows that $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - f^* \leq 2\epsilon$ for $k \geq \bar{k}$. ■

Lemma 5 (Lagrangian of Averages)**.** *Consider optimisation problem P and update $\boldsymbol{\lambda}_{k+1} = [\boldsymbol{\lambda}_k + \alpha \boldsymbol{g}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]}$ where $\alpha > 0$ and $\{\boldsymbol{z}_k\}$ is a sequence of points from C such that $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - q(\boldsymbol{\lambda}_k) \leq 2\epsilon$ for all $k = 1, 2, \ldots$. Let $\lambda_1^{(j)} \in [0, \bar{\lambda}]$ where $\bar{\lambda} \geq \lambda^{*(j)}$, $j = 1, \ldots, m$. Then,*

$$|L(\boldsymbol{z}_k^\diamond, \boldsymbol{\lambda}_k^\diamond) - f^*| \leq 2\epsilon + \frac{\alpha}{2} m\bar{g}^2 + \frac{m\bar{\lambda}^2}{\alpha k} \tag{8}$$

where $\boldsymbol{z}_k^\diamond := \frac{1}{k}\sum_{i=1}^k \boldsymbol{z}_{i+1}$, $\boldsymbol{\lambda}_k^\diamond := \frac{1}{k}\sum_{i=1}^k \boldsymbol{\lambda}_i$ and $\bar{g} := \max_{\boldsymbol{z} \in C} \|\boldsymbol{g}(\boldsymbol{z})\|_\infty$.

Proof: See Appendix. ■

Note that by selecting α sufficiently small in Lemma 5 we can obtain a sequence $\{\boldsymbol{\lambda}_k\}$ that changes sufficiently slowly so to satisfy the conditions of Lemma 4. Further, by Lemma 4 we can construct a sequence of primal variables that satisfy the conditions of Lemma 5 for $k \geq \bar{k}$ and it then follows that (8) is satisfied.

Lemma 5 requires that $\lambda^{*(j)} \leq \bar{\lambda}$ for all $j = 1, \ldots, m$, so it naturally arises the question as to when $\lambda^{*(j)}$ (and so $\bar{\lambda}$) is bounded. This is clarified in the next lemma, which corresponds to Lemma 1 in [20].

Lemma 6 (Bounded Multipliers)**.** *Let $Q_\delta := \{\boldsymbol{\lambda} \succeq \boldsymbol{0} : q(\boldsymbol{\lambda}) \geq q(\boldsymbol{\lambda}^*) - \delta\}$ with $\delta \geq 0$ and let the Slater condition hold,* i.e., *there exists a vector $\bar{\boldsymbol{z}} \in C$ such that $\boldsymbol{g}(\bar{\boldsymbol{z}}) \prec \boldsymbol{0}$. Then, for every $\boldsymbol{\lambda} \in Q_\delta$ we have that*

$$\|\boldsymbol{\lambda}\|_2 \leq \frac{1}{\upsilon}(f(\bar{\boldsymbol{z}}) - q(\boldsymbol{\lambda}^*) + \delta) \tag{9}$$

where $\upsilon := \min_{j \in \{1, \ldots, m\}} -g^{(j)}(\bar{\boldsymbol{z}})$.

Proof: First of all recall that since the Slater condition holds we have strong duality, *i.e.,* $q(\boldsymbol{\lambda}^*) = f^*$, and f^* is finite by Proposition 2.1.1. in [21]. Now observe that when $\boldsymbol{\lambda} \in Q_\delta$ then $q(\boldsymbol{\lambda}^*) - \delta \leq q(\boldsymbol{\lambda}) = \min_{\boldsymbol{z} \in C} L(\boldsymbol{z}, \boldsymbol{\lambda}) \leq f(\bar{\boldsymbol{z}}) + \boldsymbol{\lambda}^T \boldsymbol{g}(\bar{\boldsymbol{z}})$, and rearranging terms we obtain $-\boldsymbol{\lambda}^T \boldsymbol{g}(\bar{\boldsymbol{z}}) = -\sum_{j=1}^m \lambda^{(j)} g^{(j)}(\bar{\boldsymbol{z}}) \leq f(\bar{\boldsymbol{z}}) - q(\boldsymbol{\lambda}^*) + \delta$. Next, since $\boldsymbol{\lambda} \succeq \boldsymbol{0}$ and $-g^{(j)}(\bar{\boldsymbol{z}}) > 0$ for all $j = 1, \ldots, m$, let $\upsilon :=$

$\min_{j\in\{1,\dots,m\}} -g^{(j)}(\bar{\boldsymbol{z}})$ and see that $\upsilon\sum_{j=1}^m \lambda^{(j)} \le f(\bar{\boldsymbol{z}}) - q(\boldsymbol{\lambda}^*) + \delta$. Finally, dividing by υ and using the fact that $\|\boldsymbol{\lambda}\|_2 \le \sum_{j=1}^m \lambda^{(j)}$ the stated result follows. ∎

From Lemma 6 we have that it is sufficient for C_0 to have non-empty relative interior in order for Q_δ to be a bounded subset in $\mathbb{R}^m$, and since by definition $\boldsymbol{\lambda}^* \in Q_\delta$ then $\boldsymbol{\lambda}^*$ is bounded. The bound obtained in Lemma 6 depends on $q(\boldsymbol{\lambda}^*) = f^*$, which is usually not known. Nevertheless, we can obtain a looser bound if we use the fact that $-q(\boldsymbol{\lambda}^*) \le -q(\boldsymbol{\lambda})$ for all $\boldsymbol{\lambda} \succeq \mathbf{0}$. That is, for every $\boldsymbol{\lambda} \in Q_\delta$ we have that

$$\|\boldsymbol{\lambda}\|_2 \le \frac{1}{\upsilon}(f(\bar{\boldsymbol{z}}) - q(\boldsymbol{\lambda}_0) + \delta),$$

where $\boldsymbol{\lambda}_0$ is an arbitrary vector in $\mathbb{R}^m_+$.

That is, when the Slater condition is satisfied the upper and lower bounds in (8) are finite and can be made arbitrarily small as $k \to \infty$ by selecting the step size α sufficiently small. Convergence of the average of the Lagrangians does not, of course, guarantee that $f(\boldsymbol{z}_k^\diamond) \to f^*$ unless we also have complementary slackness, *i.e.,* $(\boldsymbol{\lambda}_k^\diamond)^T \boldsymbol{g}(\boldsymbol{z}_k^\diamond) \to 0$. Next we present the following lemma, which is a generalisation of Lemma 3 in [20].

Lemma 7 (Complementary Slackness and Feasibility)**.** *Let the Slater condition hold and suppose $\{\boldsymbol{z}_k\}$ is a sequence of points in C and $\{\tilde{\boldsymbol{\lambda}}_k\}$ a sequence of points in $\mathbb{R}^m$ such that (i) $L(\boldsymbol{z}_{k+1}, \tilde{\boldsymbol{\lambda}}_k) - q(\tilde{\boldsymbol{\lambda}}_k) \le 2\epsilon$ for all k and (ii) $|\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \le \alpha\sigma_0$, $j = 1,\dots,m$ where $\boldsymbol{\lambda}_{k+1} = [\boldsymbol{\lambda}_k + \alpha \boldsymbol{g}(\boldsymbol{z}_{k+1})]^+$, $\epsilon \ge 0$, $\alpha > 0$, $\sigma_0 \ge 0$. Suppose also that $\lambda_1^{(j)} \in [0, \bar{\lambda}]$ with*

$$\bar{\lambda} \ge \tfrac{3}{\upsilon}(f(\bar{\boldsymbol{z}}) - q(\boldsymbol{\lambda}^*) + \delta) + \alpha m \bar{g}$$

where $\delta := \alpha(m\bar{g}^2/2 + m^2\sigma_0\bar{g}) + 2\epsilon$, $\bar{g} := \max_{\boldsymbol{z}\in C} \|\boldsymbol{g}(\boldsymbol{z})\|_\infty$, $\bar{\boldsymbol{z}}$ a Slater vector and $\upsilon := \min_{j\in\{1,\dots,m\}} -g^{(j)}(\bar{\boldsymbol{z}})$. Then, $\lambda_k^{(j)} \le \bar{\lambda}$ for all $k = 1, 2, \dots,$

$$-\frac{m\bar{\lambda}^2}{2\alpha k} - \frac{\alpha}{2} m\bar{g}^2 \le (\boldsymbol{\lambda}_k^\diamond)^T \boldsymbol{g}(\boldsymbol{z}_k^\diamond) \le \frac{m\bar{\lambda}^2}{\alpha k} \tag{10}$$

and

$$g^{(j)}(\boldsymbol{z}_k^\diamond) \le \frac{\bar{\lambda}}{\alpha k} \tag{11}$$

where $\boldsymbol{z}_k^\diamond := \frac{1}{k}\sum_{i=1}^k \boldsymbol{z}_{i+1}$ and $\boldsymbol{\lambda}_k^\diamond := \frac{1}{k}\sum_{i=1}^k \boldsymbol{\lambda}_i$.

Proof: See Appendix. ∎

Lemma 7 is expressed in a general form where $\tilde{\boldsymbol{\lambda}}$ may be any suitable approximation to the usual Lagrange multiplier. Evidently, the lemma also applies in the special case where $\boldsymbol{\lambda}_k = \tilde{\boldsymbol{\lambda}}_k$ in which case $\sigma_0 = 0$. Note from the lemma as well that the running average $\boldsymbol{z}_k^\diamond$ is asymptotically attracted to the feasible region as k increases, *i.e.,* $\lim_{k\to\infty} \boldsymbol{g}(\boldsymbol{z}_k^\diamond) \preceq \mathbf{0}$

We are now in a position to present one of our main results:

Theorem 3 (Constrained Optimisation)**.** *Consider constrained convex optimisation P and the associated sequence of non-convex optimisations $\{\tilde{P}_k\}$:*

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} \; L((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}, \tilde{\boldsymbol{\lambda}}_k) \tag{12}$$

$$\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k \tag{13}$$

$$\boldsymbol{\lambda}_{k+1} = [\boldsymbol{\lambda}_k + \alpha \boldsymbol{g}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]} \tag{14}$$

Let the Slater condition hold and suppose that $|\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \le \alpha\sigma_0$ for all $j = 1,\dots,m$, $k \ge 1$ with $\sigma_0 \ge 0$. Further, suppose parameters α and β are selected sufficiently small that

$$0 < \alpha \le \gamma_1\gamma\beta\epsilon/(m^2(\bar{g}^2 + 2\sigma_0\bar{g})) \tag{15}$$

$$0 < \beta \le (1-\gamma)\min\{\epsilon/(\mu_L \bar{x}_D^2), 1\} \tag{16}$$

with $\epsilon > 0$, $\gamma \in (0,1)$, $\gamma_1 \in (0, 1/2)$, $\bar{x}_D := 2\max_{\boldsymbol{x}\in D}\|\boldsymbol{x}\|_2$, $\mu_L = \mu_f + \bar{\lambda}\mathbf{1}^T\boldsymbol{\mu}_{\boldsymbol{g}}$ and $\bar{\lambda}$ as given in Lemma 7. Then, for every $\epsilon > 0$ and for k sufficiently large ($k \ge \bar{k}$) the sequence of solutions $\{\boldsymbol{z}_k\}$ to sequence of optimisations $\{\tilde{P}_k\}$ satisfies:

$$-\frac{2m\bar{\lambda}^2}{\alpha k} - \alpha(m\bar{g}^2/2 + m^2\sigma_0\bar{g}) - 2\epsilon \le f(\boldsymbol{z}_k^\diamond) - f^* \le 2\epsilon + \alpha(m\bar{g}^2 + m^2\sigma_0\bar{g}) + \frac{3m\bar{\lambda}^2}{2\alpha k} \tag{17}$$

where $\boldsymbol{z}_k^\diamond := \frac{1}{k}\sum_{i=\bar{k}}^{\bar{k}+k} \boldsymbol{z}_{i+1}$, $\tilde{\boldsymbol{\lambda}}_k^\diamond := \frac{1}{k}\sum_{i=\bar{k}}^{\bar{k}+k} \tilde{\boldsymbol{\lambda}}_i$ and $\bar{g} := \max_{\boldsymbol{z}\in C}\|\boldsymbol{g}(\boldsymbol{z})\|_\infty$.

Proof: First of all observe that since $\lambda_{k+1}^{(j)} = [\lambda_k^{(j)} + \alpha g^{(j)}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]}$ we have that $|\lambda_{k+1}^{(j)} - \lambda_k^{(j)}| \le \alpha\bar{g}$ for all k. Further, since $|\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \le \alpha\sigma_0$ then $|\tilde{\lambda}_{k+1}^{(j)} - \tilde{\lambda}_k^{(j)}| = |\tilde{\lambda}_{k+1}^{(j)} - \tilde{\lambda}_k^{(j)} + \lambda_{k+1}^{(j)} - \lambda_{k+1}^{(j)} + \lambda_k^{(j)} - \lambda_k^{(j)}| \le |\tilde{\lambda}_{k+1}^{(j)} - \lambda_{k+1}^{(j)}| + |\lambda_{k+1}^{(j)} - \lambda_k^{(j)}| + |\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \le \alpha(2\sigma_0 + \bar{g})$. That is,

$$\|\tilde{\boldsymbol{\lambda}}_{k+1} - \tilde{\boldsymbol{\lambda}}_k\|_2 \le \alpha m(2\sigma_0 + \bar{g}) \qquad k = 1, 2, \dots \tag{18}$$

Next, see that since $L(\cdot, \boldsymbol{\lambda}_k)$ has uniform bounded curvature and $|L(\boldsymbol{z}, \tilde{\boldsymbol{\lambda}}_{k+1}) - L(\boldsymbol{z}, \tilde{\boldsymbol{\lambda}}_k)| \le \|\tilde{\boldsymbol{\lambda}}_{k+1} - \tilde{\boldsymbol{\lambda}}_k\|_2 \|\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2 \le \|\tilde{\boldsymbol{\lambda}}_{k+1} - \tilde{\boldsymbol{\lambda}}_k\|_2 m\bar{g} \le \alpha m^2(2\sigma_0\bar{g} + \bar{g}^2) \le \gamma_1\gamma\beta\epsilon$, it follows by Lemma 4 that for k sufficiently large ($k \ge \bar{k}$) then $L(\boldsymbol{z}_{k+1}, \tilde{\boldsymbol{\lambda}}_k) - q(\tilde{\boldsymbol{\lambda}}_k) \le 2\epsilon$ and therefore by Lemma 5

$$-\frac{m\bar{\lambda}^2}{\alpha k} - \frac{\alpha}{2} m\bar{g}^2 - 2\epsilon \le L(\boldsymbol{z}_k^\diamond, \tilde{\boldsymbol{\lambda}}_k^\diamond) - f^* \le 2\epsilon + \frac{\alpha}{2} m\bar{g}^2 + \frac{m\bar{\lambda}^2}{\alpha k}.$$

Next, see that since $|L(\boldsymbol{z}_k^\diamond, \boldsymbol{\lambda}_k^\diamond) - L(\boldsymbol{z}_k^\diamond, \tilde{\boldsymbol{\lambda}}_k^\diamond)| = (\boldsymbol{\lambda}_k^\diamond - \tilde{\boldsymbol{\lambda}}_k^\diamond)^T \boldsymbol{g}(\boldsymbol{z}_k^\diamond) \le \|\boldsymbol{\lambda}_k^\diamond - \tilde{\boldsymbol{\lambda}}_k^\diamond\|_2 \|\boldsymbol{g}(\boldsymbol{z}_k^\diamond)\|_2 \le \alpha m^2\sigma_0\bar{g}$ we have that

$$-\frac{m\bar{\lambda}^2}{\alpha k} - \alpha(m\bar{g}^2/2 + m^2\sigma_0\bar{g}) - 2\epsilon \le L(\boldsymbol{z}_k^\diamond, \boldsymbol{\lambda}_k^\diamond) - f^* \le 2\epsilon + \alpha(m\bar{g}^2/2 + m^2\sigma_0\bar{g}) + \frac{m\bar{\lambda}^2}{\alpha k}.$$

Finally, by using the complementary slackness bound of Lemma 7 the stated result follows. ∎

Theorem 3 says that by selecting step size α and smoothing parameter β sufficiently small then the average of the solutions to the sequence of non-convex optimisations $\{\tilde{P}_k\}$ can be made arbitrarily close to the solution of constrained convex optimisation P.

1) Alternative Update: Note that, by replacing use of Theorem 1 by Theorem 2 in the proof, we can replace update (12) by its non-convex Frank-Wolfe alternative,

$$\begin{aligned}\boldsymbol{x}_k &\in \arg\min_{\boldsymbol{x}\in D} \; \partial_{\boldsymbol{z}} L(\boldsymbol{z}_k, \boldsymbol{\lambda}_k)^T \boldsymbol{x} \\ &= \arg\min_{\boldsymbol{x}\in D} \; (\partial f(\boldsymbol{z}_k) + \boldsymbol{\lambda}_k^T \partial \boldsymbol{g}(\boldsymbol{z}_k))^T \boldsymbol{x}.\end{aligned} \tag{19}$$

That is, we have:

Corollary 2 (Constrained Optimisation Using Frank-Wolfe Update)**.** *Consider the setup in Theorem 3 but with update (12) replaced by (19). Then, there exists a finite $\bar{k}$ such that the bound given in (17) holds.*

C. Generalised Update

Let $C' \subseteq \mathrm{conv}(D)$ be any subset of the convex hull of action set D, including the empty set. Since $\min_{\boldsymbol{x}\in C'\cup D}\ L((1-\beta)\boldsymbol{z}_k+\beta\boldsymbol{x},\boldsymbol{\lambda}_k) \leq \min_{\boldsymbol{x}\in D}\ L((1-\beta)\boldsymbol{z}_k+\beta\boldsymbol{x},\boldsymbol{\lambda}_k)$, we can immediately generalise update (12) to

$$\boldsymbol{x}_k \in \operatorname*{arg\,min}_{\boldsymbol{x}\in C'\cup D} L((1-\beta)\boldsymbol{z}_k+\beta\boldsymbol{x},\boldsymbol{\lambda}_k) \tag{20}$$

and Theorem 3 will continue to apply. Selecting C' equal to the empty set we recover (12) as a special case. Selecting $C' = \mathrm{conv}(D)$ we recover the classical convex dual subgradient update as a special case. Update (20) therefore naturally generalises both the classical convex dual subgradient update and non-convex update (12). Hence, we have the following corollary.

Corollary 3 (Constrained Optimisation Using Unified Update)**.** *Consider the setup in Theorem 3 but with update (12) replaced by (20). Then, there exists a finite $\bar{k}$ such that the bound given in (17) holds.*

V. Using Queues As Approximate Multipliers

In Theorem 3 the only requirement on the sequence of approximate multipliers $\{\tilde{\boldsymbol{\lambda}}_k\}$ is that it remains close to the sequence of Lagrange multipliers $\{\boldsymbol{\lambda}_k\}$ generated by a dual subgradient update in the sense that $|\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \leq \alpha\sigma_0$ for all k. In this section we consider the special case where sequence $\{\tilde{\boldsymbol{\lambda}}_k\}$ additionally satisfies the following,

$$\tilde{\boldsymbol{\lambda}}_{k+1} = [\tilde{\boldsymbol{\lambda}}_k + \tilde{\boldsymbol{\delta}}_k]^{[0,\bar{\lambda}]} \tag{21}$$

with $\tilde{\boldsymbol{\delta}}_k \in \mathbb{R}^m$ and $\tilde{\boldsymbol{\lambda}}_1 = \boldsymbol{\lambda}_1$.

We begin by recalling the following lemma, which is a direct result of [22, Proposition 3.1.2].

Lemma 8. *Consider sequences $\{\lambda_k\}$, $\{\tilde{\lambda}_k\}$ in $\mathbb{R}$ given by updates $\lambda_{k+1} = [\lambda_k + \delta_k]^{[0,\bar{\lambda}]}$, $\tilde{\lambda}_{k+1} = [\tilde{\lambda}_k + \tilde{\delta}_k]^{[0,\bar{\lambda}]}$ where $\delta, \tilde{\delta} \in \mathbb{R}$. Suppose $\lambda_1 = \tilde{\lambda}_1$ and $|\sum_{i=1}^k \delta_i - \tilde{\delta}_i| \leq \epsilon$ for all k. Then,*

$$|\lambda_k - \tilde{\lambda}_k| \leq 2\epsilon \qquad k = 1, 2, \ldots$$

Proof: See Appendix. ∎

Applying Lemma 8 to our present context it follows that $|\lambda_k^{(j)} - \tilde{\lambda}_k^{(j)}| \leq \alpha\sigma_0$ for all k (and so Theorem 3 holds) for every sequence $\{\tilde{\boldsymbol{\delta}}_k\}$ such that $|\sum_{i=1}^k \alpha g^{(j)}(\boldsymbol{z}_i) - \tilde{\delta}_i^{(j)}| \leq \alpha\sigma_0$ for all k.

Of particular interest is the special case of optimisation P where the constraints are linear. That is, $g^{(j)}(\boldsymbol{z}) = \boldsymbol{a}^{(j)}\boldsymbol{z} - b^{(j)}$ where $(\boldsymbol{a}^{(j)})^T \in \mathbb{R}^n$ and $b^{(j)} \in \mathbb{R}$, $j = 1, \ldots, m$. Gathering vectors $\boldsymbol{a}^{(j)}$ together as the rows of matrix $\boldsymbol{A} \in \mathbb{R}^{m\times n}$ and collecting additive terms $b^{(j)}$ into vector $\boldsymbol{b} \in \mathbb{R}^m$, the linear constraints can then be written as $\boldsymbol{A}\boldsymbol{z} \preceq \boldsymbol{b}$. Therefore, the dual subgradient Lagrange multiplier update in the sequence of optimisations $\{\tilde{P}_k\}$ is given by

$$\boldsymbol{\lambda}_{k+1} = [\boldsymbol{\lambda}_k + \alpha(\boldsymbol{A}\boldsymbol{z}_{k+1} - \boldsymbol{b})]^{[0,\bar{\lambda}]} \tag{22}$$

with $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$, $\boldsymbol{x}_k \in D$. Now suppose that in (21) we select $\tilde{\boldsymbol{\delta}}_k = \alpha(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k)$ where $\{\boldsymbol{b}_k\}$ is a sequence of points in $\mathbb{R}^m$. Then,

$$\tilde{\boldsymbol{\lambda}}_{k+1} = [\tilde{\boldsymbol{\lambda}}_k + \alpha(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k)]^{[0,\bar{\lambda}]} \tag{23}$$

with $\tilde{\boldsymbol{\lambda}}_1 = \boldsymbol{\lambda}_1$.

Observe that in (23) we have replaced the continuous-valued quantity $\boldsymbol{z}_k$ with the discrete-valued quantity $\boldsymbol{x}_k$. We have also replaced the constant $\boldsymbol{b}$ with the time-varying quantity $\boldsymbol{b}_k$. Further, letting $\boldsymbol{Q} := \tilde{\boldsymbol{\lambda}}/\alpha$ then (23) can be rewritten equivalently as

$$\boldsymbol{Q}_{k+1} = [\boldsymbol{Q}_k + \boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k]^{[0,\bar{\lambda}/\alpha]} \tag{24}$$

which is a discrete queue length update with increment $\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k$. The approximate multipliers $\tilde{\boldsymbol{\lambda}}$ are therefore scaled discrete queue occupancies.

Using Lemma 8 it follows immediately that Theorem 3 holds provided

$$|\textstyle\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_i - \boldsymbol{x}_i) + (b_i^{(j)} - b^{(j)})| \leq \alpha\sigma_0 \tag{25}$$

Since update $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$ yields a running average of $\{\boldsymbol{x}_k\}$ we might expect that sequences $\{\boldsymbol{z}_k\}$ and $\{\boldsymbol{x}_k\}$ are always close and so uniform boundedness of $|\sum_{i=1}^k (b_i^{(j)} - b^{(j)})|$ is sufficient to ensure that (25) is satisfied. This is indeed the case, as established by the following theorem.

Theorem 4 (Queues as approximate multipliers)**.** *Consider updates (22) and (23) where $\{\boldsymbol{x}_k\}$ is an arbitrary sequence of points in D, $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$, $\beta \in (0,1)$, $\boldsymbol{z}_1 \in C := \mathrm{conv}(D)$. Further, suppose that $\{\boldsymbol{b}_k\}$ is a sequence of points in $\mathbb{R}^m$ such that $|\sum_{i=1}^k (b_i^{(j)} - b^{(j)})| \leq \sigma_2$ for all $j = 1, \ldots, m$, $k = 1, 2, \ldots$. Then,*

$$\|\tilde{\boldsymbol{\lambda}}_k - \boldsymbol{\lambda}_k\|_2 \leq 2m\alpha(\sigma_1/\beta + \sigma_2), \quad k = 1, 2, \ldots$$

where $\sigma_1 := 2\max_{\boldsymbol{z}\in C} \|\boldsymbol{A}\boldsymbol{z}\|_\infty$.

Proof: See Appendix. ∎

Observe that the difference between $\boldsymbol{\lambda}_k$ and $\tilde{\boldsymbol{\lambda}}_k$ can be made arbitrarily small by selecting α small enough. The requirement that $|\sum_{i=1}^k (b_i^{(j)} - b^{(j)})| \leq \sigma_2$ is satisfied when sequence $\{b_k^{(j)}\}$ converges sufficiently fast to $b^{(j)}$ (dividing both sides by k, the requirement is that $|\frac{1}{k}\sum_{i=1}^k b_i^{(j)} - b^{(j)}| \leq \sigma_2/k$).

In the special case when $b_k^{(j)} = b^{(j)}$ then Theorem 4 is trivially satisfied. This is illustrated in Figure 3, which plots $\boldsymbol{\lambda}_k$ and $\tilde{\boldsymbol{\lambda}}_k$ for a simple example where $A = 1$, $b_k = b = 0.5$, $\alpha = 1$, $\beta = 0.1$ and sequence $\{x_k\}$ takes values independently and uniformly at random from set $\{0, 1\}$. It can be seen that the distance between $\boldsymbol{\lambda}_k$ and $\tilde{\boldsymbol{\lambda}}_k$ remains uniformly bounded over time.

In summary, we have arrived at the following corollary to Theorem 3.

Corollary 4 (Constrained Optimisation Using Approximate Multipliers)**.** *Consider the setup of Theorem 3, suppose the*

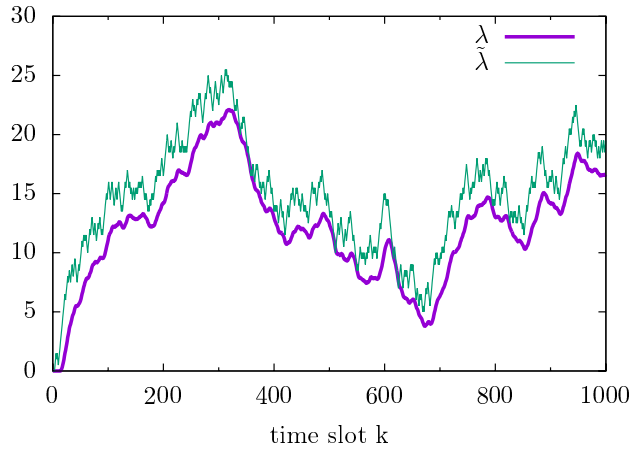

Fig. 3: Example realisations of $\tilde{\lambda}_k$ (thin line) and λ_k (thicker line) given by updates (22) and (23).

constraints are linear $\boldsymbol{A}\boldsymbol{z}-\boldsymbol{b} \preceq \boldsymbol{0}$ *and* $\tilde{\boldsymbol{\lambda}}_{k+1} = [\tilde{\boldsymbol{\lambda}}_k + \alpha(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k)]^{[0,\bar{\lambda}]}$, $\boldsymbol{b}_k \in \mathbb{R}^m$. *Suppose* $|\frac{1}{k}\sum_{i=1}^k b_i^{(j)} - b^{(j)}| \leq \sigma_2/k$ *for all* j *and* k. *Then, the bound (17) in Theorem 3 holds with* $\sigma_0 = 2(\sigma_1/\beta + \sigma_2)$ *where* $\sigma_1 := 2\max_{\boldsymbol{z}\in C} \|\boldsymbol{A}\boldsymbol{z}\|_\infty$.

A. Weaker Condition for Loose Constraints

Suppose constraint j is loose at an optimum, *i.e.*, $g^{(j)}(\boldsymbol{z}^*) < 0$ for $\boldsymbol{z}^* \in C^*$. Then by complementary slackness the associated Lagrange multiplier must be zero, *i.e.*, $(\lambda^{(j)})^* = 0$, and we can select $\lambda_k^{(j)} = (\lambda^{(j)})^* = 0$ for all k. Since $\tilde{\lambda}_k^{(j)}$ is non-negative, to apply Theorem 3 it is enough that $\tilde{\lambda}_k^{(j)} \leq \alpha\sigma_0$ for $k = 1, 2, \ldots$. Assuming, for simplicity, that $\tilde{\lambda}_1^{(j)} = 0$, from the proof of Lemma 8 we have $\tilde{\lambda}_k^{(j)} = [\max_{1\leq l\leq k-1}\sum_{i=l}^{k-1}\alpha(\boldsymbol{a}^{(j)}\boldsymbol{x}_i - b_i^{(j)})]^+$ and so a sufficient condition for $\tilde{\lambda}_k^{(j)} \leq \alpha\sigma_0$ is that $\max_{1\leq l\leq k-1}\sum_{i=l}^{k-1}(\boldsymbol{a}^{(j)}\boldsymbol{x}_i - b^{(j)}) - (b_i^{(j)} - b^{(j)}) \leq \sigma_0$ for all k. The advantage of this condition is that $-\sum_{i=l}^{k-1}(b_i^{(j)} - b^{(j)})$ is now not bounded below and so a wider class of sequences $\{b_i^{(j)}\}$ is potentially admissible. The disadvantage is that to exploit this we need to know in advance that constraint j is loose at the optimum.

B. Queue Stability

Recall that by Lemma 7 sequence $\{\boldsymbol{\lambda}_k\}$ in Theorem 3 (and respective corollaries of the theorem) is bounded for all $k \geq \bar{k}$. Therefore, since $\|\tilde{\boldsymbol{\lambda}}_k - \boldsymbol{\lambda}_k\|_2$ is uniformly bounded it follows that $\{\tilde{\boldsymbol{\lambda}}_k\}$ is also bounded and therefore the associated discrete queue is stable (although the occupancy $\boldsymbol{Q}$ of the discrete queue scales with $1/\alpha$ since $\boldsymbol{Q} = \tilde{\boldsymbol{\lambda}}/\alpha$). Note that we have arrived to this queue stability result purely from a convex optimisation analysis and without using any Foster-Lyapunov argument.

C. Optimal Actions Depend Only on Queue Occupancy

In network resource allocation problems where the linear constraints can be identified with link queues we can use the scaled queue occupancies directly in the optimisation. That is,

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} L((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}, \alpha\boldsymbol{Q}_k) \tag{26}$$

$$= \arg\min_{\boldsymbol{x}\in D} f((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}) + \alpha\beta\boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x} \tag{27}$$

where update (27) is obtained from (26) by retaining only the parts of $L((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}, \tilde{\boldsymbol{\lambda}}_k)$ which depend on $\boldsymbol{x}$ *i.e.*, dropping constant terms which do not change the solution to the optimisation. We could also consider Corollary 2 and so have a Frank-Wolfe like update:

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} \partial_{\boldsymbol{z}} L(\boldsymbol{z}_k, \alpha\boldsymbol{Q}_k)^T\boldsymbol{x} \tag{28}$$

$$= \arg\min_{\boldsymbol{x}\in D} \partial f(\boldsymbol{z}_k)^T\boldsymbol{x} + \alpha\boldsymbol{Q}_k^T\boldsymbol{A}\boldsymbol{x} \tag{29}$$

Importantly, note that neither (27) nor (29) involve $\boldsymbol{b}$ or $\boldsymbol{b}_k$. Therefore, we can generate a sequence of discrete actions by simply looking at the queue occupancies at each time slot.

VI. Stochastic Optimisation

The analysis in the preceeding Section IV is for deterministic optimisation problems. However, it can be readily extended to a class of stochastic optimisations.

A. Stochastic Approximate Multipliers

1) Linear Constraints: Of particular interest, in view of the equivalence which has been established between approximate multipliers and queues, is accommodating stochastic queue arrivals.

Let $\{\boldsymbol{B}_k\}$ be a stochastic process with realisations of $\boldsymbol{B}_k$ taking values in $\mathbb{R}^m$ and with mean $\boldsymbol{b}$. Let $p_K := \text{Prob}(\max_{k\in\{1,2,\ldots,K\}} \|\sum_{i=1}^k(\boldsymbol{B}_i - \boldsymbol{b})\|_\infty \leq \sigma_2)$. Let $\{\boldsymbol{b}_i\}_{i=1}^K$ denote a realisation of length K. Fraction p_K of these realisations satisfy $\|\sum_{i=1}^k(\boldsymbol{b}_i - \boldsymbol{b})\|_\infty \leq \sigma_2$ for all $k = 1, 2, \ldots, K$. When this fraction is asymptotically lower bounded $\liminf_{K\to\infty} p_K \geq p$, then fraction p of realisations satisfy the conditions of Theorem 4. We therefore have the following corollary (which is a stochastic version of Corollary 4) to Theorem 3.

Corollary 5. *Consider the setup of Theorem 3, suppose the constraints are linear* $\boldsymbol{A}\boldsymbol{z} - \boldsymbol{b} \preceq \boldsymbol{0}$ *and* $\tilde{\boldsymbol{\lambda}}_{k+1} = [\tilde{\boldsymbol{\lambda}}_k + \alpha(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}_k)]^{[0,\bar{\lambda}]}$, $\boldsymbol{b}_k \in \mathbb{R}^m$, $\tilde{\boldsymbol{\lambda}}_1 = \boldsymbol{\lambda}_1$. *Let sequence* $\{\boldsymbol{b}_k\}$ *be a realisation of a stochastic process* $\{\boldsymbol{B}_k\}$ *and* $p_K := \text{Prob}(\max_{k\in\{1,2,\ldots,K\}} \|\sum_{i=1}^k(\boldsymbol{B}_i - \boldsymbol{b})\|_\infty \leq \sigma_2)$. *Suppose that this probability is asymptotically lower bounded* $\liminf_{K\to\infty} p_K \geq p$. *Then, with probability at least* p *the bound (17) in Theorem 3 holds with* $\sigma_0 = 2(\sigma_1/\beta + \sigma_2)$, $\sigma_1 := 2\max_{\boldsymbol{z}\in C} \|\boldsymbol{A}\boldsymbol{z}\|_\infty$.

Note that there is no requirement for stochastic process $\{\boldsymbol{B}_k\}$ to be i.i.d. or for any of its properties, other than that feasible set $\boldsymbol{A}\boldsymbol{z} \preceq \boldsymbol{b}$ has non-empty relative interior, to be known in advance in order to construct solution sequence $\{\tilde{P}_k\}$.

2) Nonlinear Constraints: We can further generalise the latter corollary to consider non-linear stochastic constraints:

Corollary 6. *Consider the setup in Theorem 3 with update* $\tilde{\boldsymbol{\lambda}}_{k+1} = [\tilde{\boldsymbol{\lambda}}_k + \alpha(\boldsymbol{g}(\boldsymbol{z}_{k+1}) - \boldsymbol{b}_k)]^{[0,\bar{\lambda}]}$, $\boldsymbol{b}_k \in \mathbb{R}^m$, $\tilde{\boldsymbol{\lambda}}_1 = \boldsymbol{\lambda}_1$. *Let sequence* $\{\boldsymbol{b}_k\}$ *is a realisation of a stochastic process* $\{\boldsymbol{B}_k\}$ *with* $\mathbf{0}$ *mean and* $p_K := \mathrm{Prob}(\max_{k\in\{1,2,\ldots,K\}} \| \sum_{i=1}^k (\boldsymbol{B}_i - \boldsymbol{b})\|_\infty \leq \sigma_2)$. *Suppose that this probability is asymptotically lower bounded* $\liminf_{K\to\infty} p_K \geq p$. *Then, with probability at least* p *the bound (17) in Theorem 3 holds with* $\sigma_0 = 2\sigma_2$.

B. Stochastic Actions

Suppose that when at time k we select action $\boldsymbol{x}_k \in D$, the action actually applied is a realisation of random variable $\boldsymbol{Y}_k$ that also takes values in D; this is for simplicity, the extension to random action sets different from D is straightforward. For example, we may select $x_k = 1$ (which might correspond to transmitting a packet) but with some probability actually apply $y_k = 0$ (which might correspond to a transmission failure/packet loss). Let $p_{\boldsymbol{xy}} := \mathrm{Prob}(\boldsymbol{Y}_k = \boldsymbol{y} | \boldsymbol{x}_k = \boldsymbol{x})$, $\boldsymbol{x}$, $\boldsymbol{y} \in D$ and we assume that this probability distribution is time-invariant *i.e.,* does not depend on k; again, this can be relaxed in the obvious manner.

Namely, assume that the probabilities $p_{\boldsymbol{xy}}$, $\boldsymbol{x}$, $\boldsymbol{y} \in D$ are known. Then $\bar{\boldsymbol{y}}(\boldsymbol{x}) := \mathbb{E}[\boldsymbol{Y}_k | \boldsymbol{x}_k = \boldsymbol{x}] = \sum_{\boldsymbol{y}\in D} \boldsymbol{y} p_{\boldsymbol{xy}}$ can be calculated. The above analysis now carries over unchanged provided we modify the non-convex optimisation from $\min_{\boldsymbol{x}\in D} L((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}, \boldsymbol{\lambda}_k)$ to $\min_{\boldsymbol{x}\in D} L((1-\beta)\boldsymbol{z}_k + \beta\bar{\boldsymbol{y}}(\boldsymbol{x}), \boldsymbol{\lambda}_k)$ and everywhere replace $\boldsymbol{x}_k$ by $\bar{\boldsymbol{y}}(\boldsymbol{x}_k)$. That is, we simply change variables to $\bar{\boldsymbol{y}}$. Note that this relies upon the mapping from $\boldsymbol{x}$ to $\bar{\boldsymbol{y}}$ being known. If this is not the case, then we are entering the realm of stochastic decision problems and we leave this to future work.

VII. Max-Weight Revisited

A. Discussion

Recall the formulation of a queueing network in Section I, where matrix $\boldsymbol{A}$ defines the queue interconnection, with i'th row having a -1 at the i'th entry, 1 at entries corresponding to queues from which packets are sent to queue i, and 0 entries elsewhere. Hence, the queue occupancy evolves as $\boldsymbol{Q}_{k+1} = [\boldsymbol{Q}_k + \boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{b}_k]^{[0,\bar{\lambda}/\alpha]}$. As shown in Section V updates $\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} \partial f(\boldsymbol{z}_k)^T \boldsymbol{x} + \alpha \boldsymbol{Q}_k^T \boldsymbol{A}\boldsymbol{x}$, $\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k$ leads to $\boldsymbol{z}_k$ converging to a ball around the solution to the following convex optimisation,

$$\begin{array}{ll} \underset{\boldsymbol{z}\in C}{\text{minimise}} & f(\boldsymbol{z}) \\ \text{subject to} & \boldsymbol{A}\boldsymbol{z} + \boldsymbol{b} \preceq 0 \end{array}$$

where $C = \mathrm{conv}(D)$, $\{\boldsymbol{b}_k\}$ is any sequence such that $\lim_{k\to\infty} \frac{1}{k}\sum_{i=1}^k \boldsymbol{b}_i = \boldsymbol{b}$ and $|(\frac{1}{k}\sum_{i=1}^k b_i^{(j)}) - b^{(j)}| \leq \sigma_2/k$, $j = 1, \ldots, m$ for some finite $\sigma_2 > 0$.

Observe that this update is identical to the greedy primal-dual max-weight schedule once we identify utility function $U(\cdot)$ with $-f(\cdot)$. However, we have arrived at this from a purely convex optimisation viewpoint and by elementary arguments, without recourse to more sophisticated Lyapunov drift, stochastic queueing theory *etc*. Further, our analysis immediately generalises the max-weight analysis to allow arbitrary linear constraints rather than just the specific constraints associated with a queueing network, and beyond this to convex nonlinear constraints with bounded curvature.

In our analysis, the key role played by bounded curvature in non-convex descent is brought to the fore. This property is of course present in existing max-weight results, in the form of a requirement for continuous differentiability of the utility function, but insight into the fundamental nature of this requirement had been lacking. One immediate benefit is the resulting observation that any non-convex update with suitable descent properties can be used, and strong connections are established with the wealth of convex descent methods. For example, by Theorem 3 we can replace update $\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} (\partial f(\boldsymbol{z}_k) + \boldsymbol{A}\boldsymbol{Q}_k)^T \boldsymbol{x}$ (which is now seen to be a variant of the classical Frank-Wolfe update) with the direct Lagrangian update $\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} f(\boldsymbol{z}_k + \beta(\boldsymbol{x} - \boldsymbol{z}_k)) + \beta \boldsymbol{Q}_k^T \boldsymbol{A}\boldsymbol{x}$ to obtain a new class of non-convex algorithms.

VIII. Numerical Examples

A. Example: Convergence and Bounds in Theorem 3

Consider the convex optimisation problem $\min_{\boldsymbol{z}\in C} \sum_{i=1}^n \exp(\boldsymbol{V}\boldsymbol{z})$ s.t. $\boldsymbol{b} \preceq \boldsymbol{z}$ where $\boldsymbol{V} := \mathrm{diag}(1, \ldots, n)$, $C := \mathrm{conv}(D)$, $D := \{\boldsymbol{x} \in \mathbb{R}^n : x^{(j)} \in \{0, s\}, j = 1, \ldots, n\}$, $s > 0$ and $\boldsymbol{b} = (s/\sum_{i=1}^n 2i)[1, \ldots, n]^T$. Observe that the Slater condition holds. Consider the following sequence of non-convex optimisations for $k = 1, 2, \ldots,$

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{x}\in D} \sum_{i=1}^n \exp(\boldsymbol{V}((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x})) + \tilde{\boldsymbol{\lambda}}_k^T(\boldsymbol{b} - ((1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x})) \quad (30)$$

$$\boldsymbol{z}_{k+1} = (1-\beta)\boldsymbol{z}_k + \beta\boldsymbol{x}_k \quad (31)$$

$$\boldsymbol{\lambda}_{k+1} = [\boldsymbol{\lambda}_k + \alpha(\boldsymbol{b} - \boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]} \quad (32)$$

with $\boldsymbol{z}_1 = s\mathbf{1}$, $\lambda_1^{(j)}, \tilde{\lambda}_1^{(j)} = 0$, $j = 1, \ldots, m$ and parameters α and β are selected as indicated in (15) and (16), with parameters $n = 3$, $s = 1/\sqrt{n\bar{\mu}_L}$, $\bar{\mu}_L = 0.6$, $\gamma = 0.5$, $\bar{\lambda} = 0.7$, $\bar{g} = 0.6211$ and $\bar{x}_D = s\sqrt{n}$.

1) Convergence into 2ϵ*-ball in finite time:* To begin with, suppose $\tilde{\lambda} = \lambda$. For $\epsilon = 0.05$ (so $\alpha = 7.29 \cdot 10^{-5}$) Figure 4 plots the convergence of $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k)$ into an 2ϵ-ball around f^*. It can be seen that this convergence occurs within finite time, $\bar{k} = 81$ and that $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k)$ then stays within this ball at times $k \geq \bar{k}$.

2) Upper and lower bounds from Theorem 3: Now suppose that $\tilde{\lambda}^{(j)} = \lambda^{(j)} + \alpha Y_k \sigma_0$ where Y_k is uniformly randomly distributed between -1 and 1. For $\sigma_0 \in \{0, 1, 4\}$ (so $\alpha \in \{7.29 \cdot 10^{-5}, 1.85 \cdot 10^{-5}, 5.74 \cdot 10^{-6}\}$), Figure 5 plots $f(\boldsymbol{z}_k^\diamond)$ and the upper and lower bounds from Theorem 3 vs k. Figure 6 shows detail from Figure 5. It can be seen that, as expected, $f(\boldsymbol{z}_k^\diamond)$ is indeed upper and lower bounded by the values from Theorem 3. It can also be seen that the upper and lower bounds are not tight, but they are not excessively loose either.

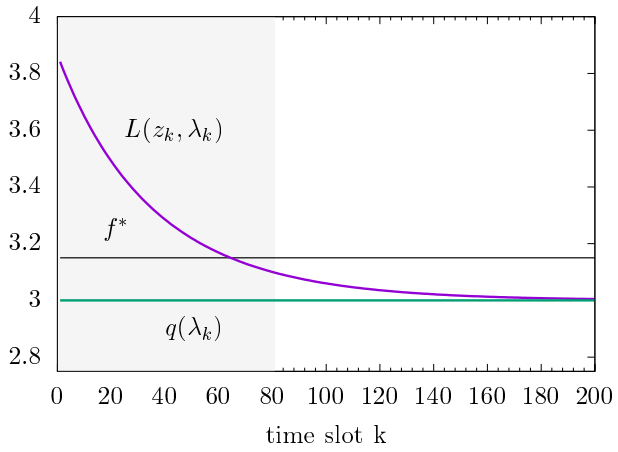

Fig. 4: Illustrating the convergence of $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k)$ to a ball around $q(\boldsymbol{\lambda}_k)$ for $\epsilon = 0.05$ and $\sigma_0 = 0$. Shaded area ($k < 81$) indicates that $L(\boldsymbol{z}_{k+1}, \boldsymbol{\lambda}_k) - q(\boldsymbol{\lambda}_k) > 2\epsilon$.

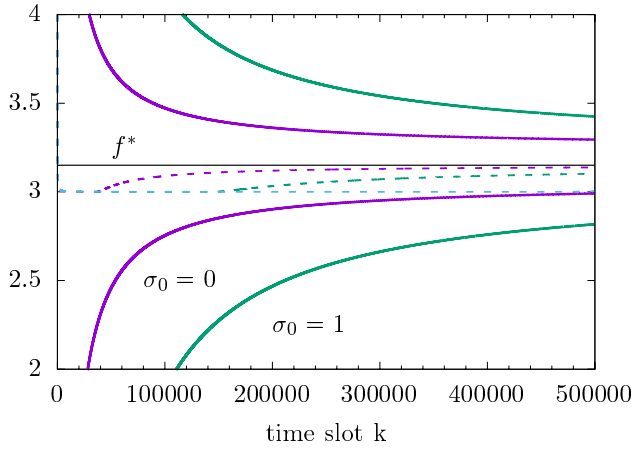

Fig. 5: Illustrating the convergence of $f(\boldsymbol{z}_k^\diamond)$ to a ball around f^* (straight line) of Example VIII-A when $\epsilon = 0.05$ and $\sigma_0 \in \{0, 1, 4\}$. Dashed lines indicate $f(\boldsymbol{z}_k^\diamond)$ with $\bar{k} = 81$ while thick lines indicate upper and lower bounds of Theorem 3.

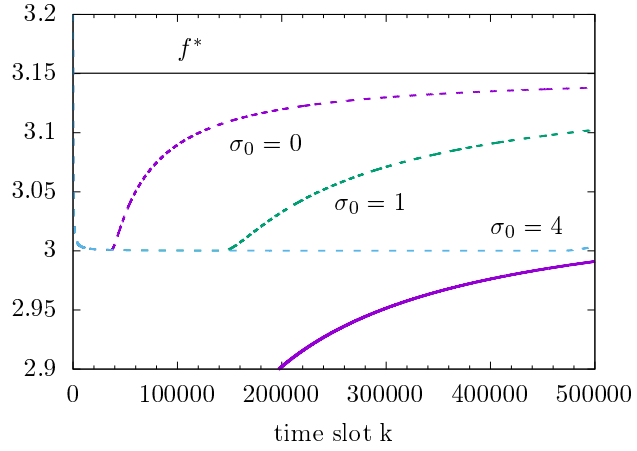

Fig. 6: Detail from Figure 5.

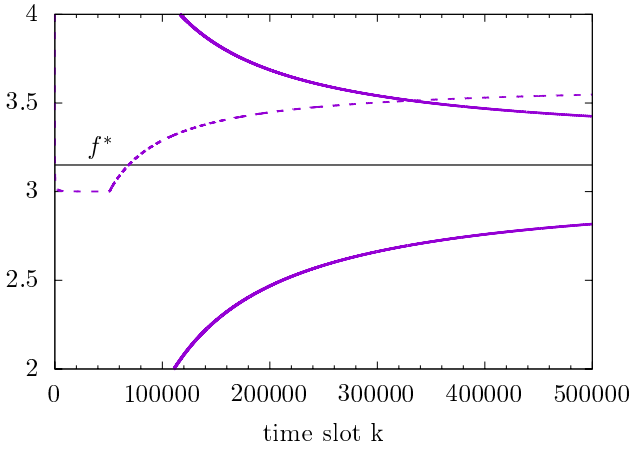

Fig. 7: Illustrating the violation of the bounds of Theorem 3 when $\tilde{\lambda}_k^{(j)} = [\lambda_k^{(j)} + \alpha e^k e^{-10^5}]^{[0,\bar{\lambda}]}$. Dashed line indicates $f(\boldsymbol{z}_k^\diamond)$, $\bar{k} = 81$, while thicker lines indicate upper and lower bounds around f^* (straight line).

3) Violation of upper bound: Let $\tilde{\lambda}_k^{(j)} = [\lambda_k^{(j)} + \alpha e^k e^{-10^5}]^{[0,\bar{\lambda}]}$. With this choice the difference between $\lambda_k^{(j)}$ and $\tilde{\lambda}_k^{(j)}$ is uniformly bounded by $\alpha\sigma_0$ with $\sigma_0 = 1$ for $k \leq 10^5$ but after that increases exponentially with k. Figure 7 plots $f(\boldsymbol{z}_k^\diamond)$ and the upper and lower bounds from Theorem 3 when parameter α is selected according to Theorem 3 assuming $\sigma_0 = 1$. It can be seen that the upper and lower bounds hold for $k \leq 10^5$, but as the difference between multipliers increases $f(\boldsymbol{z}_k^\diamond)$ is not attracted to f^* and it ends up violating the bounds.

B. Example: Privacy-Enhancing Online Scheduling

We now present a small but interesting application example which illustrates some of the generality of Theorem 3.

Consider a sequence of information packets indexed by $k = 1, 2, \ldots$. Time is slotted and the packets arrive at a queue with inter-arrival times $\{b_k\}$, $k = 1, 2, \ldots$ *i.e.*, $b_k \in \mathbb{N}$ is the number of slots between the arrival of packet $k-1$ and packet k, with $b_1 := 0$. Outgoing packet j is forwarded with inter-service time $s_j \in \{0, 1, \ldots, T\} \subset \mathbb{N}$ *i.e.*, with s_j slots between packet x_j and the previously transmitted packet. Dummy packets are transmitted as needed when no information packets are available, so as allow s_j to be freely selected and to prevent large inter-arrivals times from propagating to the outgoing packet stream. The aim is to select the queue service x_j such that the entropy of the inter-packet times of the outgoing packet stream is at least E, in order to provide a degree of resistence to traffic timing analysis, while stabilising the queue.

The packet arrival process is not known in advance, other than the facts that it can be feasibily served, the inter-arrival times have finite mean $\lim_{k\to\infty} \frac{1}{k}\sum_{i=1}^k b_i = b$ and $|(\frac{1}{k}\sum_{i=1}^k b_i) - b| \leq \sigma_2/k$ for some finite $\sigma_2 > 0$.

Suppose the inter-service times s_j are i.i.d. and let vector $\boldsymbol{p}$ with elements $p^{(i)} = \text{Prob}(s_j = i)$, $i = 0, \ldots, T$ describe the probability distribution over set $\{0, 1, \ldots, T\}$. The task can

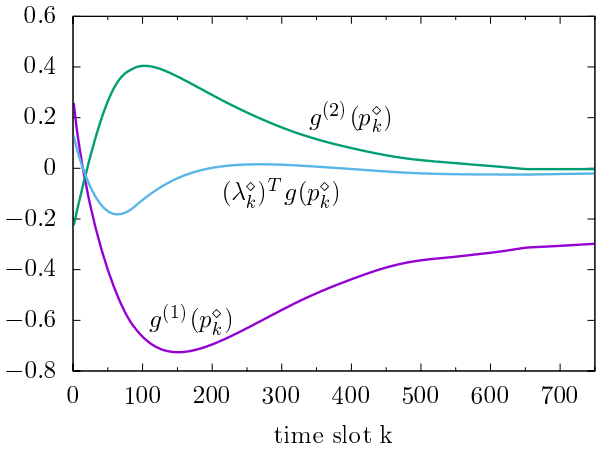

Fig. 8: Illustrating the convergence of sequence $\{\boldsymbol{p}_k\}$ into the feasible region. Updates (33) - (36) use parameters $T_{\max} = 5$, $E = \log(5)/5$, sequence $\{b_k\} = \{0, 1, 0, 1, 0, \dots\}$ so $b = 1/2$, $\xi = b/2$, $\lambda_1^{(1)} = \bar{\lambda}$, $\lambda_1^{(2)} = 0$, $\bar{\lambda} = 1/2$ and $\alpha = \beta = 0.01$.

be formulated as the following feasibility problem (couched in convex optimisation form),

$$\min_{\boldsymbol{p}\in C} 1 \text{ s.t. } \sum_{i=0}^{T_{\max}} p^{(i)} \log p^{(i)} \leq -E, \ \sum_{i=0}^{T_{\max}} i p^{(i)} + \xi \leq b$$

where $\xi > 0$ ensures that the mean inter-service time is strictly less than the mean inter-arrival time, so ensuring queue stability, and $C := \{\boldsymbol{p} \in [0,1]^T : \sum_{i=1}^{T} p^{(i)} \leq 1\}$. If the arrival process $\{b_k\}$ were known in advance, we could solve this optimisation to determine a feasible $\boldsymbol{p}$. When the arrivals are not known in advance, using generalised update (20) by Corollary 3 we can instead use the following online update to determine a sequence $\{\boldsymbol{p}_k\}$ that converges to a feasible point.

$$\boldsymbol{x}_k \in \arg\min_{\boldsymbol{p}\in C} \lambda_k^{(1)} g^{(1)}(\boldsymbol{p}) + \lambda_k^{(2)} g^{(2)}(\boldsymbol{p}) \tag{33}$$

$$\boldsymbol{p}_{k+1} = (1-\beta)\boldsymbol{p}_k + \beta \boldsymbol{x}_k \tag{34}$$

$$\lambda_{k+1}^{(1)} = \left[\lambda_k^{(1)} + \alpha g^{(1)}(\boldsymbol{p}_{k+1})\right]^{[0,\bar{\lambda}]} \tag{35}$$

$$\lambda_{k+1}^{(2)} = \left[\lambda_k^{(2)} + \alpha (g^{(2)}(\boldsymbol{p}_{k+1}) + b - b_k)\right]^{[0,\bar{\lambda}]} \tag{36}$$

where $g^{(1)}(\boldsymbol{p}) := \sum_{i=0}^{T_{\max}} p^{(i)} \log p^{(i)} + E$ and $g^{(2)}(\boldsymbol{p}) := \sum_{i=0}^{T_{\max}} i p^{(i)} + \xi - b$ with $\lambda_1^{(1)}, \lambda_1^{(2)} \in [0, \bar{\lambda}]$. The online update does not require knowledge of the mean inter-arrival time b since in (33) the $\arg\min$ does not depend on b while in (36) we have $g^{(2)}(\boldsymbol{p}_{k+1}) + b - b_k = \sum_{i=0}^{T_{\max}} i p_{k+1}^{(i)} + \xi - b_k$.

Figure 8 illustrates the online update. It can be seen that approximate complementary slackness converges to a ball around 0, and that constraints $g^{(j)}(\boldsymbol{p}_k^\diamond),\ j = 1, 2$ are attracted to the feasible region as k increases.

We highlight the following aspects of this example:

1) The online update differs from the standard dual-subgradient ascent in its use of the observed inter-arrival times b_k rather than the (unknown) mean inter-arrival time b. The inter-arrival times b_k are discrete-valued, which also takes us outside of the usual setting. The great advantage of the online update is that does not require knowledge of the mean rate b of the packet arrival process, which is unknown beforehand, but only makes myopic use of available measurements to construct a packet schedule.

2) The constraint $\sum_{i=0}^{T_{\max}} i p^{(i)} < b$ is expressed in terms of the packet inter-arrival and inter-service times rather than the number of packet arrivals and departures. Hence, θ_k is not the scaled link queue occupancy but rather is related to the scaled queue waiting time. Note that θ_k is not exactly the waiting time since the mean value $\sum_{i=0}^{T_{\max}} i p^{(i)}$ is used for the inter-service time rather than the actual inter-service time realisations.

3) The transmission of dummy packets is explicitly an aspect of the application and it is the *transmitted* packets (both dummy and information packets) which matter for satisfying the entropy constraint, not just the information packets. This is because it is packet timing rather than packet content which is of interest here.

4) Decision variable x_k is not a packet or commodity scheduling action and so there are no issues around not having a packet to send when a queue is empty.

5) The entropy constraint is highly nonlinear and not at all like the type of flow constraint encountered in typical queueing network applications.

IX. Conclusions

In this paper we investigate the connections between max-weight approaches and dual subgradient methods for convex optimisation. We find that strong connections do indeed exist and we establish a clean, unifying theoretical framework that includes both max-weight and dual subgradient approaches as special cases.

X. Appendix: Proofs

A. Proof of Theorems 1 and 2

The following two fundamental results are the key to establishing Theorem 1:

Lemma 9. *Let $D := \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_{|D|}\}$ be a finite set of points from $\mathbb{R}^n$ and $C := \text{conv}(D)$. Then, for any point $\boldsymbol{y} \in C$ and vector $\boldsymbol{z} \in \mathbb{R}^n$ there exists a point $\boldsymbol{x} \in D$ such that $\boldsymbol{z}^T(\boldsymbol{x} - \boldsymbol{y}) \leq 0$.*

Proof: Since $\boldsymbol{y} \in C := \text{conv}(D)$, $\boldsymbol{y} = \sum_{j=1}^{|D|} \theta^{(j)} \boldsymbol{x}_j$ with $\sum_{j=1}^{|D|} \theta^{(j)} = 1$ and $\theta^{(j)} \in [0,1]$. Hence, $\boldsymbol{z}^T(\boldsymbol{x} - \boldsymbol{y}) = \sum_{j=1}^{|D|} \theta^{(j)} \boldsymbol{z}^T(\boldsymbol{x} - \boldsymbol{x}_j)$. Select $\boldsymbol{x} \in \arg\min_{\boldsymbol{w}\in D} \boldsymbol{z}^T \boldsymbol{w}$. Then $\boldsymbol{z}^T\boldsymbol{x} \leq \boldsymbol{z}^T\boldsymbol{x}_j$ for all $\boldsymbol{x}_j \in D$ and so $\boldsymbol{z}^T(\boldsymbol{x} - \boldsymbol{y}) \leq 0$. ■

Lemma 10 (Non-Convex Descent)**.** *Let $F(\boldsymbol{z})$ be a convex function and suppose points $\boldsymbol{y}, \boldsymbol{z} \in C = \text{conv}(D)$ exist such that $F(\boldsymbol{y}) \leq F(\boldsymbol{z}) - \epsilon$, $\epsilon > 0$. Suppose $F(\cdot)$ has bounded curvature on C with curvature constant μ_F. Then there exists at least one $\boldsymbol{x} \in D$ such that $F((1-\beta)\boldsymbol{z} + \beta\boldsymbol{x}) \leq F(\boldsymbol{z}) - \gamma\beta\epsilon$ with $\gamma \in (0,1)$ provided $0 < \beta \leq (1-\gamma)\min\{\epsilon/(\mu_F \bar{x}_D^2), 1\}$.*

Proof: By convexity,

$$F(\boldsymbol{z}) + \partial F(\boldsymbol{z})^T(\boldsymbol{y} - \boldsymbol{z}) \leq F(\boldsymbol{y}) \leq F(\boldsymbol{z}) - \epsilon.$$

Hence, $\partial F(\boldsymbol{z})^T(\boldsymbol{y}-\boldsymbol{z}) \leq -\epsilon$. Now observe that for $\boldsymbol{x} \in D$ we have $(1-\beta)\boldsymbol{z}+\beta\boldsymbol{x} \in C$ and by the bounded curvature of $F(\cdot)$ on C

$$\begin{aligned}
&F((1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}) \\
&\leq F(\boldsymbol{z}) + \beta\partial F(\boldsymbol{z})^T(\boldsymbol{x}-\boldsymbol{z}) + \mu_F\beta^2\|\boldsymbol{x}-\boldsymbol{z}\|_2^2 \\
&= F(\boldsymbol{z}) + \beta\partial F(\boldsymbol{z})^T(\boldsymbol{y}-\boldsymbol{z}) + \beta\partial F(\boldsymbol{z})^T(\boldsymbol{x}-\boldsymbol{y}) + \mu_F\beta^2\|\boldsymbol{x}-\boldsymbol{z}\|_2^2 \\
&\leq F(\boldsymbol{z}) - \beta\epsilon + \beta\partial F(\boldsymbol{z})^T(\boldsymbol{x}-\boldsymbol{y}) + \mu_F\beta^2\|\boldsymbol{x}-\boldsymbol{z}\|_2^2
\end{aligned}$$

By Lemma 9 we can select $\boldsymbol{x} \in D$ such that $\partial F(\boldsymbol{z})^T(\boldsymbol{x}-\boldsymbol{y}) \leq 0$. With this choice of $\boldsymbol{x}$ it follows that

$$\begin{aligned}
F((1-\beta)\boldsymbol{z}+\beta\boldsymbol{x}) &\leq F(\boldsymbol{z}) - \beta\epsilon + \mu_F\beta^2\|\boldsymbol{x}-\boldsymbol{z}\|_2^2 \\
&\leq F(\boldsymbol{z}) - \beta\epsilon + \mu_F\beta^2\bar{x}_D^2 \qquad (37)
\end{aligned}$$

where (37) follows from Lemma 2, and the result now follows. ∎

Proof of Theorem 1: Since $F_k(\cdot)$ has bounded curvature for any k it is continuous, and as C is closed and bounded we have by the Weierstrass theorem (*e.g.*, see Proposition 2.1.1 in [21]) that $\min_{\boldsymbol{z}\in C} F_k(\boldsymbol{z})$ is finite. We now proceed considering two cases:

Case (i): $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) \geq \epsilon$. By Lemma 10 there exists $\boldsymbol{x}_k \in D$ such that $F_k((1-\beta)\boldsymbol{z}_k+\beta\boldsymbol{x}_k) - F_k(\boldsymbol{z}_k) = F_k(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{z}_k) \leq -\gamma\beta\epsilon$. Further, since $F_{k+1}(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{z}_{k+1}) \leq \gamma_1\gamma\beta\epsilon$ and $F_k(\boldsymbol{z}_k) - F_{k+1}(\boldsymbol{z}_k) \leq \gamma_1\gamma\beta\epsilon$ it follows

$$F_{k+1}(\boldsymbol{z}_{k+1}) - F_{k+1}(\boldsymbol{z}_k) \leq 2\gamma_1\gamma\beta\epsilon - \gamma\beta\epsilon < 0. \qquad (38)$$

That is, $F_k(\cdot)$ and $F_{k+1}(\cdot)$ decrease monotonically when $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) \geq \epsilon$.

Case (ii): $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) < \epsilon$. It follows that $F_k(\boldsymbol{z}_k) < F_k(\boldsymbol{y}_k^*)+\epsilon$. Since $F_k(\cdot)$ is convex and has bounded curvature, $F_k(\boldsymbol{z}_{k+1}) \leq F_k(\boldsymbol{z}_k) + \beta\partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) + \beta^2\mu_F\bar{x}_D^2 \leq F_k(\boldsymbol{y}_k^*) + \epsilon + \beta\partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) + \beta^2\mu_F\bar{x}_D^2$. The final term holds uniformly for all $\boldsymbol{x}_k \in D$ and since we select $\boldsymbol{x}_k$ to minimise $F_k(\boldsymbol{z}_{k+1})$ by Lemma 9 we therefore have $F_k(\boldsymbol{z}_{k+1}) \leq F_k(\boldsymbol{y}_k^*) + \epsilon + \beta^2\mu_F\bar{x}_D^2$. Using the stated choice of β and the fact that $F_{k+1}(\boldsymbol{z}_{k+1}) - \gamma_1\gamma\beta\epsilon \leq F_k(\boldsymbol{z}_{k+1})$ yields

$$F_{k+1}(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{y}_k^*) \leq \gamma_1\gamma\beta\epsilon + \epsilon + \beta(1-\gamma)\epsilon. \qquad (39)$$

Finally, since $F_k(\boldsymbol{y}_k^*) \leq F_k(\boldsymbol{y}_{k+1}^*) \leq F_{k+1}(\boldsymbol{y}_{k+1}^*) + \gamma_1\gamma\beta\epsilon$ we obtain

$$\begin{aligned}
F_{k+1}(\boldsymbol{z}_{k+1}) - F_{k+1}(\boldsymbol{y}_{k+1}^*) &\leq 2\gamma_1\gamma\beta\epsilon + \epsilon + \beta(1-\gamma)\epsilon \\
&\leq 2\epsilon.
\end{aligned}$$

We therefore have that $F_{k+1}(\boldsymbol{z}_k)$ is strictly decreasing when $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) \geq \epsilon$ and otherwise uniformly upper bounded by 2ϵ. It follows that for all k sufficiently large $F_k(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{y}_k^*) \leq 2\epsilon$ as claimed. ∎

Proof of Theorem 2: Firstly, we make the following observations,

$$\begin{aligned}
&\arg\min_{\boldsymbol{z}\in C} F_k(\boldsymbol{z}_k) + \partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{z}-\boldsymbol{z}_k) \\
&\overset{(a)}{=} \arg\min_{\boldsymbol{z}\in C} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{z} \overset{(b)}{=} \arg\min_{\boldsymbol{x}\in D} \partial F_k(\boldsymbol{z}_k)^T\boldsymbol{x} \qquad (40)
\end{aligned}$$

where equality (a) follows by dropping terms not involving $\boldsymbol{z}$ and (b) from the observation that we have a linear program (the objective is linear and set C is a polytope, so defined by linear constraints) and so the optimum lies at an extreme point of set C *i.e.,* in set D. We also have that

$$\begin{aligned}
F_k(\boldsymbol{z}_k) + \partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) &\overset{(a)}{\leq} F_k(\boldsymbol{z}_k) + \partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{y}_k^*-\boldsymbol{z}_k) \\
&\overset{(b)}{\leq} F_k(\boldsymbol{y}_k^*) \leq F_k(\boldsymbol{z}_k)
\end{aligned}$$

where $\boldsymbol{y}_k^* \in \arg\min_{\boldsymbol{z}\in C} F_k(\boldsymbol{z})$, inequality (a) follows from the minimality of $\boldsymbol{x}_k$ in C noted above and (b) from the convexity of $F_k(\cdot)$. It follows that $\partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) \leq -(F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*)) \leq 0$. We have two cases to consider. Case (i): $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) \geq \epsilon$. By the bounded curvature of $F_k(\cdot)$,

$$\begin{aligned}
F_k(\boldsymbol{z}_{k+1}) &\leq F_k(\boldsymbol{z}_k) + \beta\partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) + \mu_F\beta^2\bar{x}_D \\
&\leq F_k(\boldsymbol{z}_k) - \beta\epsilon + \mu_f\beta^2\bar{x}_D \leq F_k(\boldsymbol{z}_k) - \gamma\beta\epsilon.
\end{aligned}$$

Hence,

$$\begin{aligned}
F_{k+1}(\boldsymbol{z}_{k+1}) &\leq F_k(\boldsymbol{z}_{k+1}) + |F_{k+1}(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{z}_{k+1})| \\
&\leq F_k(\boldsymbol{z}_k) - \gamma\beta\epsilon + \gamma_1\gamma\beta\epsilon,
\end{aligned}$$

and since $F_k(\boldsymbol{z}_k) \leq F_{k+1}(\boldsymbol{z}_k) + \gamma_1\gamma\beta\epsilon$ we have that $F_{k+1}(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{z}_{k+1}) < 0$. Case (ii): $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) < \epsilon$. Then

$$\begin{aligned}
F_k(\boldsymbol{z}_{k+1}) &\leq F_k(\boldsymbol{z}_k) + \beta\partial F_k(\boldsymbol{z}_k)^T(\boldsymbol{x}_k-\boldsymbol{z}_k) + \mu_F\beta^2\bar{x}_D \\
&\leq F_k(\boldsymbol{y}_k^*) + \epsilon + \beta\epsilon,
\end{aligned}$$

and similar to the proof of Theorem 1 we obtain that $F_{k+1}(\boldsymbol{z}_{k+1}) - F_{k+1}(\boldsymbol{y}_{k+1}^*) \leq 2\epsilon$. We therefore have that $F_k(\boldsymbol{z}_k)$ is strictly decreasing when $F_k(\boldsymbol{z}_k) - F_k(\boldsymbol{y}_k^*) \geq \epsilon$ and otherwise uniformly upper bounded by 2ϵ. Thus for k sufficiently large $F_k(\boldsymbol{z}_{k+1}) - F_k(\boldsymbol{y}_k^*) \leq 2\epsilon$. ∎

Proof of Lemma 5: Let $\boldsymbol{\theta} \in \mathbb{R}^m$ such that $\theta^{(j)} \leq \bar{\lambda}$ for all $j = 1, \ldots, m$ and see that

$$\begin{aligned}
&\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\theta}\|_2^2 \\
&= \|[\boldsymbol{\lambda}_k + \alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]} - \boldsymbol{\theta}\|_2^2 \\
&\leq \|[\boldsymbol{\lambda}_k + \alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^{+} - \boldsymbol{\theta}\|_2^2 \qquad (41) \\
&\leq \|\boldsymbol{\lambda}_k + \alpha\boldsymbol{g}(\boldsymbol{z}_{k+1}) - \boldsymbol{\theta}\|_2^2 \\
&= \|\boldsymbol{\lambda}_k - \boldsymbol{\theta}\|_2^2 + 2\alpha(\boldsymbol{\lambda}_k - \boldsymbol{\theta})^T\boldsymbol{g}(\boldsymbol{z}_{k+1}) + \alpha^2\|\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2^2 \\
&\leq \|\boldsymbol{\lambda}_k - \boldsymbol{\theta}\|_2^2 + 2\alpha(\boldsymbol{\lambda}_k - \boldsymbol{\theta})^T\boldsymbol{g}(\boldsymbol{z}_{k+1}) + \alpha^2 m\bar{g}^2, \qquad (42)
\end{aligned}$$

where (41) follows since $\bar{\lambda} \geq \theta^{(j)}$ and (42) from the fact that $\|\boldsymbol{g}(\boldsymbol{z})\|_2^2 \leq m\bar{g}^2$ for all $\boldsymbol{z} \in C$. Applying the latter argument recursively for $i = 1, \ldots, k$ yields $\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\theta}\|_2^2 \leq \|\boldsymbol{\lambda}_1 - \boldsymbol{\theta}\|_2^2 + 2\alpha\sum_{i=1}^k(\boldsymbol{\lambda}_i - \boldsymbol{\theta})^T\boldsymbol{g}(\boldsymbol{z}_{i+1}) + \alpha^2 m\bar{g}^2 k$. Rearranging terms, dividing by $2\alpha k$, and using the fact that $\|\boldsymbol{\lambda}_{k+1} - \boldsymbol{\theta}\|_2^2 \geq 0$ and $\|\boldsymbol{\lambda}_1 - \boldsymbol{\theta}\|_2^2 \leq 2m\bar{\lambda}^2$ we have

$$\begin{aligned}
-\frac{m\bar{\lambda}^2}{\alpha k} - \frac{\alpha}{2}m\bar{g}^2 &\leq \frac{1}{k}\sum_{i=1}^k(\boldsymbol{\lambda}_i - \boldsymbol{\theta})^T\boldsymbol{g}(\boldsymbol{z}_{i+1}) \qquad (43) \\
&= \frac{1}{k}\sum_{i=1}^k L(\boldsymbol{z}_{i+1}, \boldsymbol{\lambda}_i) - L(\boldsymbol{z}_{i+1}, \boldsymbol{\theta}). \qquad (44)
\end{aligned}$$

Next, see that by the definition of sequence $\{\boldsymbol{z}_k\}$ we can write $\frac{1}{k}\sum_{i=1}^k L(\boldsymbol{z}_{i+1},\boldsymbol{\lambda}_i) \leq \frac{1}{k}\sum_{i=1}^k q(\boldsymbol{\lambda}_i)+2\epsilon \leq q(\boldsymbol{\lambda}_k^\diamond)+2\epsilon$ where the last inequality follows by the concavity of q. That is,

$$-\frac{m\bar{\lambda}^2}{\alpha k}-\frac{\alpha}{2}m\bar{g}^2-2\epsilon \leq q(\boldsymbol{\lambda}_k^\diamond)-\frac{1}{k}\sum_{i=1}^k L(\boldsymbol{z}_{i+1},\boldsymbol{\theta}) \quad (45)$$

By fixing $\boldsymbol{\theta}$ to $\boldsymbol{\lambda}^*$ and $\boldsymbol{\lambda}_k^\diamond$ and using the fact that $\frac{1}{k}\sum_{i=1}^k L(\boldsymbol{z}_{i+1},\boldsymbol{\lambda}_k^\diamond) \geq L(\boldsymbol{z}_k^\diamond,\boldsymbol{\lambda}_k^\diamond)$ for all $k=1,2,\dots$ and $\frac{1}{k}\sum_{i=1}^k L(\boldsymbol{z}_{i+1},\boldsymbol{\lambda}^*) \geq f^*$ we have that

$$-\frac{m\bar{\lambda}^2}{\alpha k}-\frac{\alpha}{2}m\bar{g}^2-2\epsilon \leq q(\boldsymbol{\lambda}_k^\diamond)-f^* \leq 0 \quad (46)$$

and

$$-\frac{m\bar{\lambda}^2}{\alpha k}-\frac{\alpha}{2}m\bar{g}^2-2\epsilon \leq q(\boldsymbol{\lambda}_k^\diamond)-L(\boldsymbol{z}_k^\diamond,\boldsymbol{\lambda}_k^\diamond) \leq 0. \quad (47)$$

Multiplying (46) by -1 and combining it with (47) yields the result. ∎

Proof of Lemma 7:

We start by showing that updates $[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+$ and $[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]}$ are interchangeable when $L(\boldsymbol{z}_{k+1},\tilde{\boldsymbol{\lambda}}_k)$ is uniformly close to $q(\tilde{\boldsymbol{\lambda}}_k)$. First of all see that

$$\begin{aligned}
&\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\|_2^2\\
&=\|[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+-\boldsymbol{\lambda}^*\|_2^2\\
&\leq\|\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})-\boldsymbol{\lambda}^*\|_2^2\\
&=\|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2^2+\alpha^2\|\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2^2+2\alpha(\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1})\\
&\leq\|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2^2+\alpha^2m\bar{g}^2+2\alpha(\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1}).
\end{aligned}$$

Now observe that since $\|\boldsymbol{\lambda}_k-\tilde{\boldsymbol{\lambda}}_k\|_2 \leq \|\boldsymbol{\lambda}_k-\tilde{\boldsymbol{\lambda}}_k\|_1 \leq \alpha m\sigma_0$ we can write $(\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1}) = (\tilde{\boldsymbol{\lambda}}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1})+(\boldsymbol{\lambda}_k-\tilde{\boldsymbol{\lambda}}_k)^T\boldsymbol{g}(\boldsymbol{z}_{k+1}) \leq (\tilde{\boldsymbol{\lambda}}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1})+\|\boldsymbol{\lambda}_k-\tilde{\boldsymbol{\lambda}}_k\|_2\|\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2 \leq (\tilde{\boldsymbol{\lambda}}_k-\boldsymbol{\lambda}^*)^T\boldsymbol{g}(\boldsymbol{z}_{k+1})+\alpha m^2\sigma_0\bar{g} = L(\boldsymbol{z}_{k+1},\tilde{\boldsymbol{\lambda}}_k)-L(\boldsymbol{z}_{k+1},\boldsymbol{\lambda}^*)+\alpha m^2\sigma_0\bar{g}$. Furthermore, since $L(\boldsymbol{z}_{k+1},\tilde{\boldsymbol{\lambda}}_k) \leq q(\tilde{\boldsymbol{\lambda}}_k)+2\epsilon$ and $-L(\boldsymbol{z}_{k+1},\boldsymbol{\lambda}^*) \leq -q(\boldsymbol{\lambda}^*)$ it follows that

$$\begin{aligned}
&\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\|_2^2-\|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2^2\\
&\quad\leq\alpha^2(m\bar{g}^2+2m^2\sigma_0\bar{g})+2\alpha(q(\tilde{\boldsymbol{\lambda}}_k)+2\epsilon-q(\boldsymbol{\lambda}^*)). \quad (48)
\end{aligned}$$

Now let $Q_\delta := \{\boldsymbol{\lambda} \succeq \mathbf{0} : q(\boldsymbol{\lambda}) \geq q(\boldsymbol{\lambda}^*)-\delta\}$ and consider two cases. Case (i) ($\tilde{\boldsymbol{\lambda}}_k \notin Q_\delta$). Then $q(\tilde{\boldsymbol{\lambda}}_k)-q(\boldsymbol{\lambda}^*) < -\delta$ and from (48) we have that $\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\|_2^2 < \|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2^2$, *i.e.,*

$$\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\|_2-\|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2<0$$

and so $\boldsymbol{\lambda}_k$ converges to a ball around $\boldsymbol{\lambda}^*$ when $\tilde{\boldsymbol{\lambda}}_k \in Q_\delta$. Case (ii) ($\tilde{\boldsymbol{\lambda}}_k \in Q_\delta$). See that $\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\| = \|[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+-\boldsymbol{\lambda}^*\|_2 \leq \|\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})-\boldsymbol{\lambda}^*\|_2 \leq \|\boldsymbol{\lambda}_k\|_2+\|\boldsymbol{\lambda}^*\|_2+\alpha m\bar{g}$. Next recall that when the Slater condition holds by Lemma 6 we have for all $\boldsymbol{\lambda} \in Q_\delta$ then $\|\boldsymbol{\lambda}\| \leq \frac{1}{v}(\eta+\delta)$ where $\eta := f(\bar{\boldsymbol{z}})-q(\boldsymbol{\lambda}^*)$ and $\bar{\boldsymbol{z}}$ a Slater vector. Therefore,

$$\|\boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}^*\|_2 \leq \frac{2}{v}(\eta+\delta)+\alpha m\bar{g}.$$

From both cases it follows that if $\|\boldsymbol{\lambda}_1-\boldsymbol{\lambda}^*\|_2 \leq \frac{2}{v}(\eta+\delta)+\alpha m\bar{g}$ then $\|\boldsymbol{\lambda}_k-\boldsymbol{\lambda}^*\|_2 \leq \frac{2}{v}(\eta+\delta)+\alpha m\bar{g}$ for all $k \geq 1$. Using this observation and the fact that $\|\boldsymbol{\lambda}_1-\boldsymbol{\lambda}^*\|_2 \geq |\|\boldsymbol{\lambda}_1\|_2-\|\boldsymbol{\lambda}^*\|_2| \geq \|\boldsymbol{\lambda}_1\|_2-\|\boldsymbol{\lambda}^*\|_2$ we obtain that when $\|\boldsymbol{\lambda}_1\|_2 \leq \frac{3}{v}(\eta+\delta)+\alpha m\bar{g}$ then $\|\boldsymbol{\lambda}_k\|_2 \leq \frac{3}{v}(\eta+\delta)+\alpha m\bar{g}$ for all $k \geq 1$. That is, if we choose $\lambda_1^{(j)} \leq \frac{3}{v}(\eta+\delta)+\alpha m\bar{g} \leq \bar{\lambda}$ then $\lambda_k^{(j)} \leq \bar{\lambda}$ for all $j=1,\dots,m$, $k \geq 1$ and so updates $[\boldsymbol{\lambda}_k+\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+$ and $[\boldsymbol{\lambda}_k+\boldsymbol{g}(\boldsymbol{z}_{k+1})]^{[0,\bar{\lambda}]}$ are interchangeable as claimed.

Now we proceed to prove the upper and lower bounds in (10). For the lower bound see first that

$$\begin{aligned}
\|\boldsymbol{\lambda}_{k+1}\|_2^2&=\|[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+\|_2^2\\
&\leq\|\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2^2\\
&=\|\boldsymbol{\lambda}_k\|_2^2+\alpha^2\|\boldsymbol{g}(\boldsymbol{z}_{k+1})\|_2^2+2\alpha\boldsymbol{\lambda}_k^T\boldsymbol{g}(\boldsymbol{z}_{k+1})\\
&\leq\|\boldsymbol{\lambda}_k\|_2^2+\alpha^2m\bar{g}^2+2\alpha\boldsymbol{\lambda}_k^T\boldsymbol{g}(\boldsymbol{z}_{k+1})
\end{aligned}$$

Rearranging terms and applying the latter bound recursively for $i=1,\dots,k$ yields $2\alpha\sum_{i=1}^k\boldsymbol{\lambda}_i^T\boldsymbol{g}(\boldsymbol{z}_{i+1}) \geq \|\boldsymbol{\lambda}_{k+1}\|_2^2-\|\boldsymbol{\lambda}_1\|_2^2-\alpha^2m\bar{g}^2k \geq -\|\boldsymbol{\lambda}_1\|_2^2-\alpha^2m\bar{g}^2k$. The bound does not depend on sequence $\{\boldsymbol{z}_k\}$, hence, it holds for any sequence of points in C. Fixing $\boldsymbol{z}_{i+1}$ to $\boldsymbol{z}_k^\diamond$ for all $i=1,\dots,k$ we can write $2\alpha\sum_{i=1}^k\boldsymbol{\lambda}_i^T\boldsymbol{g}(\boldsymbol{z}_k^\diamond) = 2\alpha k(\boldsymbol{\lambda}_k^\diamond)^T\boldsymbol{g}(\boldsymbol{z}_k^\diamond)$. Dividing by $2\alpha k$ and using the fact that $\|\boldsymbol{\lambda}_1\|_2^2 \leq m\bar{\lambda}^2$ yields

$$-\frac{m\bar{\lambda}}{2\alpha k}-\frac{\alpha}{2}m\bar{g}^2 \leq (\boldsymbol{\lambda}_k^\diamond)^T\boldsymbol{g}(\boldsymbol{z}_k^\diamond).$$

For the upper bound see that $\boldsymbol{\lambda}_{k+1}=[\boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})]^+ \succeq \boldsymbol{\lambda}_k+\alpha\boldsymbol{g}(\boldsymbol{z}_{k+1})$ and so we can write $\alpha\sum_{i=1}^k\boldsymbol{g}(\boldsymbol{z}_{i+1}) \preceq \sum_{i=1}^k(\boldsymbol{\lambda}_{i+1}-\boldsymbol{\lambda}_i) = \boldsymbol{\lambda}_{k+1}-\boldsymbol{\lambda}_1 \preceq \boldsymbol{\lambda}_{k+1}$. Next, by the convexity of $\boldsymbol{g}(\cdot)$ we have that $\frac{1}{k}\sum_{i=1}^k\alpha\boldsymbol{g}(\boldsymbol{z}_{i+1}) \succeq \alpha\boldsymbol{g}(\boldsymbol{z}_k^\diamond)$ and so it follows that $\boldsymbol{g}(\boldsymbol{z}_k^\diamond) \preceq \boldsymbol{\lambda}_{k+1}/(\alpha k)$. Multiplying the last equation by $\boldsymbol{\lambda}_k^\diamond$ and using the fact that $\mathbf{0} \preceq \boldsymbol{\lambda}_{k+1} \preceq \bar{\lambda}\mathbf{1}$ and $\mathbf{0} \preceq \boldsymbol{\lambda}_k^\diamond \preceq \bar{\lambda}\mathbf{1}$ yields the upper bound.

Finally, the constraint violation bound (11) follows from the fact that $\boldsymbol{g}(\boldsymbol{z}_k^\diamond) \preceq \bar{\lambda}^2/(\alpha k)\mathbf{1}$. ∎

Proof of Lemma 8: First of all see that $|\lambda_{k+1}-\tilde{\lambda}_{k+1}| = |[\lambda_k+\delta_k]^{[0,\bar{\lambda}]}-[\tilde{\lambda}_k+\tilde{\delta}_k]^{[0,\bar{\lambda}]}| \leq |[\lambda_k+\delta_k]^{[0,\bar{\lambda}]}-[\tilde{\lambda}_k+\tilde{\delta}_k]^+| = |[\tilde{\lambda}_k+\tilde{\delta}_k]^+-[\lambda_k+\delta_k]^{[0,\bar{\lambda}]}| \leq |[\tilde{\lambda}_k+\tilde{\delta}_k]^+-[\lambda_k+\delta_k]^+|$. We now proceed to bound the RHS of the last equation. Let $\Delta_k := -\min(\lambda_k+\delta_k,0)$, *i.e.,* $\lambda_{k+1}=\lambda_k+\delta_k+\Delta_k$ so that we can write $\lambda_{k+1}=\lambda_1+\sum_{i=1}^k(\delta_i+\Delta_i)$. Note that when $\lambda_{k+1}>0$ then $\Delta_k=0$, and that when $\lambda_{k+1}=0$ then $\sum_{i=1}^k\Delta_i=-\lambda_1-\sum_{i=1}^k\delta_i$. Next, note that since Δ_k is nonnegative for all k by construction we have that $\sum_{i=1}^k\Delta_i$ is non-decreasing in k. Using the latter observation it follows that $\sum_{i=1}^k\Delta_i=[-\lambda_1-\min_{1\leq j\leq k}\sum_{i=1}^j\delta_i]^+$ and therefore

$$\lambda_{k+1}=\sum_{i=1}^k\delta_i+\max\{\Theta_k,\lambda_1\}$$

where $\Theta_k := -\min_{1\leq j\leq k}\sum_{i=1}^j\delta_i$. Now see that

$$\begin{aligned}
&|\lambda_{k+1}-\tilde{\lambda}_{k+1}|\\
&=|\textstyle\sum_{i=1}^k\delta_i+\max\{\Theta_k,\lambda_1\}-\sum_{i=1}^k\tilde{\delta}_i-\max\{\tilde{\Theta}_k,\lambda_1\}|\\
&\leq|\textstyle\sum_{i=1}^k\delta_i-\tilde{\delta}_i|+|\max\{\Theta_k,\lambda_1\}-\max\{\tilde{\Theta}_k,\lambda_1\}|\\
&\overset{(a)}{\leq}|\textstyle\sum_{i=1}^k\delta_i-\tilde{\delta}_i|+|\tilde{\Theta}_k-\Theta_k|\\
&=|\textstyle\sum_{i=1}^k\delta_i-\tilde{\delta}_i|+|\min_{1\leq j\leq k}\sum_{i=1}^j\tilde{\delta}_j-\min_{1\leq j\leq k}\sum_{i=1}^j\delta_j|\\
&=|\textstyle\sum_{i=1}^k\delta_i-\tilde{\delta}_i|+|\max_{1\leq j\leq k}\sum_{i=1}^j-\tilde{\delta}_j-\max_{1\leq j\leq k}\sum_{i=1}^j-\delta_j|\\
&\leq|\textstyle\sum_{i=1}^k\delta_i-\tilde{\delta}_i|+\max_{1\leq j\leq k}|\sum_{i=1}^j\delta_j-\sum_{i=1}^j\tilde{\delta}_j|
\end{aligned}$$

where (a) follows easily from enumerating the four cases. Finally, since $|\sum_{i=1}^k \delta_i - \tilde{\delta}_i| \leq \max_{i\leq j\leq k} |\sum_{i=1}^j \delta_i - \tilde{\delta}_i|$ and $|\sum_{i=1}^k \delta_i - \tilde{\delta}_i| \leq \epsilon$ for all $k = 1, 2, \ldots$ the result follows. ■

Proof of Theorem 4: By Lemma 8 we require $|\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) + b_i^{(j)} - b^{(j)}|$ to be uniformly bounded in order to establish the boundedness of $|\tilde{\lambda}_k^{(j)} - \lambda_k^{(j)}|$ for all $k \geq 1$. However, since $|\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) + b_i^{(j)} - b^{(j)}| \leq |\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_{i+1} - \boldsymbol{x}_i)| + |\sum_{i=1}^k b_i^{(j)} - b^{(j)}|$ and $|\sum_{i=1}^k b_i^{(j)} - b^{(j)}| \leq \sigma_2$ by assumption, it is sufficient to show that $|\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_{i+1} - \boldsymbol{x}_i)|$ is bounded.

Now observe that since $\boldsymbol{z}_{i+1} = (1-\beta)\boldsymbol{z}_i + \beta \boldsymbol{x}_i$ we have $\boldsymbol{z}_{i+1} - \boldsymbol{x}_i = (1-\beta)(\boldsymbol{z}_i - \boldsymbol{x}_i)$. That is, $\sum_{i=1}^k (\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) = (1-\beta)\sum_{i=1}^k (\boldsymbol{z}_i - \boldsymbol{x}_i)$. Further, since $\sum_{i=1}^k (\boldsymbol{z}_i - \boldsymbol{x}_i) = \sum_{i=1}^{k-1} (\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) + (\boldsymbol{z}_1 - \boldsymbol{x}_k) = (1-\beta)\sum_{i=1}^{k-1} (\boldsymbol{z}_i - \boldsymbol{x}_i) + (\boldsymbol{z}_1 - \boldsymbol{x}_k)$ it follows that $\sum_{i=1}^k (\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) = (1-\beta)^2 \sum_{i=1}^{k-1} (\boldsymbol{z}_i - \boldsymbol{x}_i) + (1-\beta)(\boldsymbol{z}_1 - \boldsymbol{x}_k)$. Applying the preceding argument recursively we obtain that $\sum_{i=1}^k (\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) = (1-\beta)(\boldsymbol{z}_1 - \boldsymbol{x}_k) + (1-\beta)^2(\boldsymbol{z}_1 - \boldsymbol{x}_{k-1}) + \cdots + (1-\beta)^k(\boldsymbol{z}_1 - \boldsymbol{x}_1)$, *i.e.*,

$$\sum_{i=1}^k (\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) = \sum_{i=1}^k (1-\beta)^{k+1-i}(\boldsymbol{z}_1 - \boldsymbol{x}_i). \tag{49}$$

Using (49) it follows that

$$\begin{aligned} &2\alpha \left| \textstyle\sum_{i=1}^k \boldsymbol{a}^{(j)}(\boldsymbol{z}_{i+1} - \boldsymbol{x}_i) \right| \\ &\quad \leq 2\alpha \left| \textstyle\sum_{i=1}^k (1-\beta)^{k+1-i} \boldsymbol{a}^{(j)}(\boldsymbol{z}_1 - \boldsymbol{x}_i) \right| \\ &\quad \leq 2\alpha\sigma_1 \textstyle\sum_{i=1}^k (1-\beta)^{k+1-i} \end{aligned} \tag{50}$$

where $\sigma_1 := 2\max_{\boldsymbol{z}\in C} \|\boldsymbol{A}\boldsymbol{z}\|_\infty$. Next, see that $\sum_{i=1}^k (1-\beta)^{k+1-i} = (1-\beta)^{k+1}\sum_{i=1}^k (1-\beta)^{-i}$ and that

$$\sum_{i=1}^k \frac{1}{(1-\beta)^i} = \frac{1-(1-\beta)^{k+1}}{\beta(1-\beta)^k}.$$

Therefore, $\sum_{i=1}^k (1-\beta)^{-i} < (1-\beta)^{-k}/\beta$ and so

$$(1-\beta)^{k+1}\sum_{i=1}^k (1-\beta)^{-i} < \frac{(1-\beta)}{\beta} < \frac{1}{\beta}.$$

Finally, using the latter bound in (50) the stated result now follows. ■

REFERENCES

[1] A. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, no. 4, pp. 401–457, 2005.

[2] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, Dec 1992.

[3] M. Neely, E. Modiano, and C. Rohrs, "Power allocation and routing in multibeam satellites with time-varying channels," *Networking, IEEE/ACM Transactions on*, vol. 11, no. 1, pp. 138–152, Feb 2003.

[4] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 6, pp. 1333–1344, Dec 2007.

[5] A. L. Stolyar, "Greedy primal-dual algorithm for dynamic resource allocation in complex networks," *Queueing Systems*, vol. 54, no. 3, pp. 203–220, 2006.

[6] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 1, pp. 89–103, Jan 2005.

[7] M. Neely, "Energy optimal control for time-varying wireless networks," *Information Theory, IEEE Transactions on*, vol. 52, no. 7, pp. 2915–2934, July 2006.

[8] M. Neely, E. Modiano, and C. ping Li, "Fairness and optimal stochastic control for heterogeneous networks," *Networking, IEEE/ACM Transactions on*, vol. 16, no. 2, pp. 396–409, April 2008.

[9] M. Neely, *Stochastic network optimization with application to communication and queueing systems*. Morgan & Claypool Publishers, 2010.

[10] L. Georgiadis, M. Neely, and L. Tassiulas, *Resource Allocation and Cross-Layer Control in Wireless Networks*. Morgan & Claypool Publishers, 2006.

[11] M. Neely, "Distributed and secure computation of convex programs over a network of connected processors," in *Proc DCDIS Conf, Guelph, Ontario*, 2005.

[12] X. Lin, N. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Selected Areas in Communications*, vol. 24, no. 8, pp. 1452–1463, 2006.

[13] L. Huang and M. Neely, "Delay reduction via lagrange multipliers in stochastic network optimization," *IEEE Trans Automatic Control*, vol. 56, no. 4, pp. 842–857, 2011.

[14] D. P. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.

[15] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., 1989.

[16] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Research Logistics Quarterly*, vol. 3, no. 1-2, pp. 95–110, 1956.

[17] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *ICML*, 2003, pp. 928–936.

[18] A. D. Flaxman, A. T. Kalai, and H. B. McMahan, "Online convex optimization in the bandit setting: Gradient descent without a gradient," in *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '05. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2005, pp. 385–394.

[19] A. Roberts and D. Varberg, "Another proof that convex functions are locally lipschitz," *The American Mathematical Monthly*, vol. 81, no. 9, pp. 1014–1016, Nov 1974.

[20] A. Nedić and A. Ozdaglar, "Approximate primal solutions and rate analysis for dual subgradient methods," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1757–1780, 2009.

[21] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Athena Scientific, 2003.

[22] S. P. Meyn, *Control techniques for complex networks*. Cambridge University Press, 2008.